

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

NETWORK SURVIVABILITY ANALYSIS OF THE NAVY
AND MARINE CORPS INTRANET (NMCI)

by

Alex B. Fahrenthold

September 2002

Thesis Advisor:
Co-Advisor:

John Arquilla
Rex Buddenberg

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2002		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE Title Network Survivability Analysis of the Navy and Marine Corps Intranet			5. FUNDING NUMBERS	
6. AUTHOR (S) Alex B Fahrenthold			8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the U.S. Department of Defense or the U.S. Government.	
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) <p>NMCI is a mechanism to transform the Navy and Marine Corps and move them into the 21st century of warfare. Just as the Internet has transformed business and commerce around the globe, NMCI will transform the U. S. Navy and Marine Corps by harnessing the power of an integrated network. Consequently, the Navy and Marine Corps must consider systems and strategies that address the need for survivability of the network mission essential functions in the same manner applied to major weapons systems on the battlefield.</p> <p>Network Survivability is a field of study that addresses exactly this issue. Developed in 1998 under a Department of Defense contract by the Carnegie Mellon University Software Engineering Institute, Network Survivability addresses the need of a network to fulfill its essential mission in the presence of failures, compromise, or attack.</p> <p>This thesis examines the Navy and Marine Corps Intranet mission and structure in an attempt to determine its inherent survivability and ability to support the needs of the Navy and Marine Corps team. It focuses on identifying the network mission functions and the ability of the network architecture to produce the required survivability characteristics. Based on this examination I propose a mission definition for NMCI and highlight the needs within the security architecture to achieve a survivable NMCI.</p>				
14. SUBJECT TERMS NMCI, Network Survivability			15. NUMBER OF PAGES 225	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**NETWORK SURVIVABILITY ANALYSIS OF THE NAVY AND MARINE CORPS
INTRANET (NMCI)**

Alex Brian Fahrenthold
Commander, United States Naval Reserve
BBA, University of Houston, 1984

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN INFORMATION SYSTEMS TECHNOLOGY

from the

**NAVAL POSTGRADUATE SCHOOL
SEPTEMBER 2002**

Author: Alex B Fahrenthold

Approved by: John Arquilla
Thesis Advisor

Rex Buddenberg
Co-Advisor

Dan Boger
Chairman, Information Sciences Department

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

NMCI is a mechanism to transform the Navy and Marine Corps information systems and prepare 21st century warfare. Just as the Internet has transformed business and commerce around the globe, NMCI may transform the U. S. Navy and Marine Corps by harnessing the power of an integrated network. The Navy and Marine Corps Intranet constitutes the first major step into a truly network-centric warfare environment and makes them full participants in the cyber world. This network will handle the data on which an increasing percentage of the Navy and Marine Corps mission essential services will rely. Yet, the hardware and software that make up these systems have demonstrated vulnerabilities that put these mission essential functions at risk. Consequently, the Navy and Marine Corps must consider systems and strategies that address the need for survivability of the mission essential functions in the same manner applied to major weapons systems on the battlefield.

"Network survivability" is a field of study that addresses exactly this issue. Developed in 1998 under a Department of Defense contract by the Carnegie Mellon University Software Engineering Institute, network survivability addresses the need of a network to fulfill its essential mission in the presence of failures, compromise, or attack.

This thesis examines the Navy and Marine Corps Intranet mission and structure in an attempt to determine its inherent survivability and ability to support the needs

of the Navy and Marine Corps team. It focuses on identifying the network mission functions and the ability of the network architecture to produce the required survivability characteristics. Based on this examination I propose a mission definition for NMCI and highlight the need within the security architecture to achieve a survivable NMCI network.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I. NMCI OVERVIEW	1
A. WHAT IS THE NAVY AND MARINE CORPS INTRANET (NMCI)	1
B. WHY EXAMINE NMCI?	6
C. THE CONCEPT OF SURVIVABLE NETWORK SYSTEMS	11
D. NMCI SECURITY ARCHITECTURE	17
E. AVAILABILITY	25
F. QUALITY OF SERVICE	27
G. SECURITY	29
H. METHODOLOGY FOR EXAMINING NMCI	38
II. MISSION AND THE NAVY AND MARINE CORPS INTRANET	41
A. THE CURRENT MISSION DEFINITION	41
1. Alpha-Xray Battle Group Communications Example .	42
B. NMCI MISSION DEFINITION	47
C. REDEFINING THE MISSION FOR NMCI	48
D. EVALUATING NMCI MISSION CAPABILITIES	48
E. ADMINISTRATION	50
1. Availability	51
2. Security	52
3. Quality of Service	53
F. FORCE PROJECTION	53
1. Availability	54
2. Security	56
3. Quality of Service	57
G. BATTLE MANAGEMENT	60
1. Availability	61
2. Security	63
3. Quality of Service	66
H. MISSION AREA SUMMARY	67
I. THE NEW MISSION DEFINITION FOR NMCI	70
1. Mission Essential Functions	71
a. Logistics	73
b. Readiness	75
J. CONCLUSION	75
III. LEGACY SYSTEMS AND NMCI	77
A. WHAT IS LEGACY AND WHY IS IT IMPORTANT?	77
B. LEGACY TRANSITION	80
1. Legacy Networks	81
2. Legacy Applications	82
C. NMCI TRANSITION ORGANIZATIONAL METHODS	86
1. Transition by Claimant	86
2. Transition by Warfare Specialty	88

3. Transition by Navy Region	89
D. NMCI LEGACY TRANSITION	91
E. THE CURRENT PLAN FOR LEGACY TRANSITION	95
F. DESIGN METHODS FOR IMPLEMENTING LEGACY APPLICATION TRANSITION	99
G. THE CURRENT TRANSITION; GRAND DESIGN METHOD	100
H. TRANSITION UNDER THE NSA CONCEPT; SPIRAL DESIGN METHOD	103
1. The First Spiral; Roll Out (Connectivity)	105
2. The Second Spiral: Logistics	106
3. The 3rd Spiral: Readiness	108
4. Spiral Sub-Flows	109
I. CONCLUSION	111
IV. NMCI SECURITY ARCHITECTURE	115
A. INTRODUCTION	115
B. NMCI SECURITY STRATEGY	120
C. INITIATIVE CEDED TO THE ATTACKER	124
1. Intrusion Detection Systems	125
2. Firewalls	126
3. Link Encryption	127
4. Virtual Private Networks (VPN's)	127
5. Anti-Virus Software	129
6. Initiative Ceded to the Attacker: Summary	130
D. FAILURE TO ADDRESS THE MISSION OF THE NETWORK	131
E. MAHANIAN STRATEGY AND THE NETWORKED ENVIRONMENT	134
F. NMCI STRATEGY SUMMARY	136
G. NMCI SECURITY TACTICS	137
1. NMCI Availability	138
a. NMCI Availability Summary	149
2. NMCI Security	150
a. Transport Boundary	154
b. Boundary Layer 1	157
c. Boundary Layer 2	160
d. Boundary Layer 3	162
e. Boundary Layer 4	164
f. Security Summary	164
H. QUALITY OF SERVICE	169
I. NMCI SECURITY ARCHITECTURE CONCLUSION	173
V. CONCLUSION	177
A. INTRODUCTION	177
B. NMCI MISSION DEFINITION	178
C. NMCI LEGACY AND TRANSITION	179
D. NMCI SECURITY ARCHITECTURE	182
E. NMCI SURVIVABILITY CONCLUSION	184
F. RECOMMENDATIONS	184

APPENDIX A	189
APPENDIX B	199
LIST OF REFERENCES	201
INITIAL DISTRIBUTION LIST	205

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1	CERT Statistics 1990 through 2001	8
Figure 2	Classified Boundary 1	34
Figure 3	Unclassified Boundary	35
Figure 4	Legacy End to End Process	93
Figure 5	Current Transition Plan	97
Figure 6	Grand Design Method	100
Figure 7	Spiral Design Method for NMCI Implementation	104
Figure 8	The Spiral Method View of Force Projection	110
Figure 9	NMCI Seat-Centric Network Defense	121
Figure 10	Reported Cyber Incidents 1990 to 2001	133
Figure 11	NMCI Series Network Effective Availability	140

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1	Bounded vs. Unbounded Systems	12
Table 2	Survivability Characteristics	14
Table 3	Communities of Interest within NMCI	22
Table 4	ISO 7498-2 Layer Model	30
Table 5	Buddenberg Matrix of Security Requirements	31
Table 6	NMCI Misalignments within Buddenberg Matrix	33
Table 7	NMCI Contractual SLA Availability Levels	62
Table 8	Mission Requirements Summary Matrix	68
Table 9	Force Projection Mission Essential Functions	71
Table 10	SLA Availability	139
Table 11	SLA Effective Availability	142
Table 12	SLA Effective Availability vs. Large U.S. ISP	144
Table 13	Buddenberg Matrix of Security Requirements	152
Table 14	Boundary System Summary	153
Table 15	NMCI COIs	163
Table 16	ISO 7498-2 Layers	167
Table 17	Mission Area Threshold Requirements	186

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGEMENTS

To my wife Joan and my children Grace and Clark who supported me through the long days, watching an old dog begrudgingly learn new tricks.

To Dr John Arquilla for his guidance and for motivating me to take on the hard target.

To Professor Rex Buddenberg for availability, security, and quality of service; his eloquence and ability to boil the complex down to the elemental is unmatched.

To Carl von Clausewitz, for pointing the way.

THIS PAGE INTENTIONALLY LEFT BLANK

I. NMCI OVERVIEW

A. WHAT IS THE NAVY AND MARINE CORPS INTRANET (NMCI)

When beginning any examination of Navy and Marine Corps Intranet, the first consideration should be given to the definition of NMCI and what it's intended to be used for by the Navy and Marine Corps. Having a definition of NMCI will then lend context to the examination of the system as a whole and help distinguish what is relevant in the process of examination. Once we have established what NMCI is, we can begin to consider what NMCI should be.

The NMCI implementation is being overseen by the office of the Department of the Navy Chief Information Officer (DON CIO) and this office should provide the answer to the question of what NMCI is. The DON-CIO hosted website can be found at <http://www.don-imit.navy.mil>. The links provided to the program documentation found there offer a fundamental view of what the DON CIO is expecting NMCI to do for the Navy. The following is a summary of the purpose of NMCI that can be found there;

NMCI is an initiative that launches the Department of the Navy's first step toward reaching Joint Vision 2010's goal of information superiority for the Department of Defense. Defined as the ability to collect process and disseminate an uninterrupted flow of information while denying the same to an adversary, information superiority has been called the backbone of the revolution in military affairs. As DoN's first step, NMCI will establish a standardized end-to-end system for voice, video and data communications for all civilian and military personnel within the Department of the Navy. [DONIT02]

Immediately below this summary is a link that provides the Report to Congress on NMCI by the then Secretary of the Navy, the Honorable Richard Danzig, of 20 May 2000. In the first page of the Executive Summary to the Report to Congress on NMCI, Secretary Danzig explained the rationale behind the decision to deploy NMCI the following way:

The Navy Marine Corps Intranet offers the opportunity for the Department of the Navy (DON) to leverage new technologies and industry innovation to better achieve our global Naval mission. This investment in the future will build the modern Navy-Marine Corps on the transformational power of networking. It will enable connection to the National Infrastructure, extend sharing and creation of knowledge and expertise worldwide, empower innovative work and training and enhance the Quality of Life for every Marine, Sailor, and DON Civilian.

Replacing the Navy's numerous shore-based networks, NMCI will equip us with the access, interoperability, and security for our information and communications by providing voice, video and data services to all Navy and Marine Corps personnel. Coupled with the Navy's shipboard Information Technology for the 21st Century and the Marine Corps' embarked Marine Corps Tactical Network (MCTN), NMCI will provide a world-wide reach-back capability for our deployed forces.

The NMCI approach adapts what is commonly practiced in the commercial sector to acquire IT services for the government. This approach uses performance based, enterprise wide services contract that incorporates future strategic computing and communications capability and is managed much the same as a utility. [RD00]

This summary is quite effective in highlighting the three fundamental purposes that support the NMCI deployment. Taking them one at a time we can expand these statements to better understand the path that NMCI will take.

First, NMCI is a mechanism to transform the Navy and Marine Corps information systems and prepare them for 21st century warfare. Just as the Internet has transformed business and commerce around the globe, NMCI is intended to transform the U. S. Navy and Marine Corps by harnessing the power of an integrated network. NMCI is a piece of the "Global Information Grid" that is intended to support all U. S. forces deployed and in the Continental United States (CONUS) with administrative, logistical, force projection, or battlefield management data and communications [JC01]. Through the development of "virtual communities" within the Navy/Marine Corps, NMCI will leverage the Navy's knowledge base to improve how we fight, how we organize our forces, and how we manage our capital assets. NMCI is a means of moving the Navy toward network centric-warfare and transforming the organization as a whole.

Second, NMCI is a procurement strategy for the Navy's Information Technology (IT) assets. The Navy recognizes that industry, and not government, is the primary driver in IT business systems. Given its limited funding, it is logical for the DoN to follow the best practices of industry in their effort to secure IT services. NMCI shifts the burden of legacy hardware, software, and the expense of systems maintenance by procuring a service contract in the same way many large corporations have done. In doing so, the Navy maximizes its flexibility by not being tied to any single technology and being able to take advantage of new

technological advances in the market place. In the long view, NMCI is a procurement strategy that will reduce the cost of the Navy's IT service while maintaining the hardware technology of the system within one generation.

Third, NMCI is a mechanism of connectivity between all of the DoN Activities and personnel, both military and civilian. In May of 1999 the Space and Air Warfare Command (SPAWAR) identified three primary goals for the Navy and Marine Corps Intranet. They were;

- Provide quality service at a low price
- Greatly enhance information assurance of the naval enterprise
- Provide the enabler for the enterprise-wide BRP/ERP and the Revolution in Business Affairs

These goals have been refined and in the Navy and Marine Corps Intranet Brief presented to NMCI Information Bureau Oversight Council, 18 April 2001, by Mr. Joseph Cipriano, Program Executive Officer for Information Technology, identified them as follows;

- Repurposed network as a Navy-wide asset
- Bandwidth on demand
- Extend sharing and creation of knowledge and expertise worldwide
- Technology to support innovative work and training
- Make life better for every Sailor, Marine and DON Civilian

With the exception of the first two bullets from the May 1999 goals listed by SPAWAR, all of the others directly relate to the connectivity between the members of the DoN for the purpose of building community, knowledge, or

enhancing the individual quality of life. This was largely the benefit realized by both industry and the community with the advent and maturity of the Internet. Metcalf's Law is the relevant factor in this decision. Metcalf's Law states; "The value of a network grows as the square of the number of its users." [WRD98] Simply put, as more and more connections are made on a network, the more valuable each connection becomes, and the more valuable the network becomes as a whole. This is directly relevant to the implementation of NMCI and reflects the desire to achieve maximum connectivity within the DON. Looking again at the DON CIO webpage we find the following summary to describe the logic behind the procurement strategy for NMCI;

NMCI, an adaptation of what is commonly practiced in the commercial sector, represents a new approach to acquiring IT services for the government. NMCI will be a performance-based, enterprise-wide services contract that incorporates future strategic computing and communications capability and is managed much the same as any "utility." It will be purchased for the commercial sector just as we buy other types of utilities (e.g., water, telephone, gas and electricity) paying for the service as it is delivered. [RD00]

The most focused definition of NMCI provided by the DoN CIO is the function of a utility, something that provides a specific service to an end user. From this we can infer that the primary intent of NMCI is pure connectivity. If you examine the goals as they are listed from 1999 to the present, the importance of connectivity to the achievement of those goals is obvious.

If the Navy hopes to develop and foster virtual communities they must maximize connectivity throughout the Navy organization.

The same can be said for the development and sharing of knowledge bases. The connectivity provided by NMCI is the catalyst for any of this activity. High connectivity and a high availability of service are also essential for the procurement strategy to be successful. To gain acceptance and for new practices to be assimilated, any business enterprise system must provide the desired level of services to the members. If NMCI meets this objective of connectivity then the formation of community, the development and sharing of knowledge bases and improvement of the quality of life of every sailor and marine will eventually follow. For the Navy to move to the fulfillment of Network Centric Warfare and for the IT procurement strategy to succeed, connectivity must be achieved. The first two principles flow from connectivity, and so it can be argued that connectivity is the heart, if not the central purpose of NMCI and its implementation for the Navy and Marine Corps.

B. WHY EXAMINE NMCI?

The establishment of the Navy and Marine Corps Intranet (NMCI) represents a fundamental change in the business model for the U.S. Navy and Marine Corps. The deployment of NMCI will move the United States Navy into a realm not fully explored by any of her sister services, or for that matter, any other part of the United States government. While there are significant numbers of websites that represent the arms of local, state, and national government, they do not constitute full participation in the web environment by any part of the government. Web sites are a means of utilizing the cyberspace arena for one's own purpose. Fully participating in cyberspace is to assume all the risks and pursue all the possibilities and benefits that it offers,

not just the use of a portion of its capabilities for expedience or convenience. The Navy and Marine Corps Intranet (NMCI) will make the U.S. Navy a full participant in cyberspace, subject to all its potential benefits and risks. Full participation will mean the U.S. Navy will experience organizational pressures and hostile threats they have not experienced before, or even foreseen. Cyberspace, an arena now familiar to many in business and industry, is largely "Indian country" for the U.S. Navy. An examination of NMCI's preparedness for operation in this environment is a logical step given the fundamental nature of the change likely produced by NMCI's deployment.

When fully deployed, NMCI will touch every part of the Navy's organization in a way that has become a fundamental part of our business and war fighting capability, through the Navy's Information Technology Infrastructure. The business community is a good source for comparison when examining the transition of the U.S Navy into the e-business environment.

This is not meant to say that the Navy will fight battles using only electrons, but rather that more of the basic functions necessary to operate a modern armed force are and can be done via the World Wide Web. Many of the military's basic functions are already dependent upon the Internet for operation [RC02].

For example, we can look to the U.S. Army and its effort to develop a single networked enterprise. Much like a corporation, the majority of the Army's budget—60 percent—goes for salaries, business programs and systems [GCN02]. The business community is still wrestling with the effects and implications of such a fundamental connection as they work to secure their networks from

intruders that intend to do harm or steal valuable information. The e-business environment has been the proving ground, or the killing field, for intrusion detection systems (IDS), firewalls, anti-virus, and operating system software. What history has shown is that while we have been working harder to secure our networks, security failures continue to occur. Table 1-1 shows the trend of cyberspace security incidents from 1990 through 2001.




Figure 1 CERT Statistics 1990 through 2001

The trend of reported incidents was steady from 1990 through 1998, however in the 1999 through 2001 time period the number of reported incidents has more than doubled from year to year. The increases indicate that for all the effort placed on securing networks, those that wish to penetrate and damage networks are still achieving some measure of success. Firewalls, IDS's, and anti-virus software comprise an "ex post facto" defensive system. These software systems are effective at protecting the user from events that are widely known and have already occurred on a large scale, while doing little to prevent, deter, or blunt the effects of future viruses or computer network attacks. These figures also suggest that the basic approach

to network defense may be flawed. Protecting the infrastructure of national defense with "new" software systems that are dated when measured against the evolving threat may be akin to preparing for the last war.

The paradigm of security and the state of the art of network security software may be inadequate to protect NMCI and the Navy organizations, planning systems and assets it serves. Networks, especially those operated by the armed forces, need to have more than security as their framework for defending against an attack. Hardening a target will not guarantee its security or success in the face of attack. The mission-directed behavior of a network needs to be considered when designing a protective scheme [RJE99]. A network must be protected to be useful, but it must also be useful while being protected, lest it lose all relevance and value to the user. Examining alternative methods of network defense for NMCI is then reasonable, given the value of the assets being defended, the historical likelihood of an attack, and the probability that the attack would have some measure of success. An alternative method that could be applied to NMCI is Survivable Network Systems analysis, developed by the Carnegie Mellon University Software Engineering Institute (CMUSEI).

The Survivable Network Systems concept is based on the idea that a network should be designed to continue to function in the face of intrusions and compromises, and then regain full functionality soon after the intrusions or compromises end [RJE98]. Survivable Networks exhibit two essential characteristics (1) survivable networks will continue to deliver their essential services in spite of active intrusions or compromises, and (2) survivable

systems recover, in a timely manner, full mission services and capability.

Survivability is a concept readily familiar to any naval officer. Watertight compartments, a survivability feature of warships, merchant vessels, and passenger liners, have been an integral part of ship design for over 100 years. The U.S Navy's most valued and well-protected assets, the carriers, were built on the premise of survivability. With an attending force of surface and subsurface units, along with 100 combat aircraft, the designers assumed that the aircraft carriers would likely sustain combat damage. They were designed and constructed so as to be capable of absorbing significant damage while continuing to steam and operate independently.

This same logic should be applied to the Navy and Marine Corps Intranet. NMCI, it should be assumed, will be attacked and will be damaged, and like the carrier at the heart of the modern battle group, it must survive, and continue to operate and fight effectively. The examination of NMCI's survivability is therefore as relevant as evaluating the survivability of any other major weapons system deployed by the U. S. Navy and Marine Corps. The Survivable Network Systems paradigm is suitable for NMCI because it focuses on both the mission of NMCI and the protective systems of the network itself. To understand the relevance of survivability in the context of networked systems, we must understand the definition of network survivability and how it differs from the traditional approach of network security in the protection of such a significant asset as NMCI.

C. THE CONCEPT OF SURVIVABLE NETWORK SYSTEMS

To understand survivable networks we first need to define what bounded and unbounded networks are, and then describe the environment of survivable networks and where its application is relevant.

An unbounded network is a network that possesses no central administrative authority for the imposition of policy or sanctions on the members. In an unbounded system the members do not have complete visibility, execute control only within their domain, and must rely upon trust relationships among their neighbors to operate. The Internet is an example of the ultimate unbounded network. There is no central figure responsible for its content, protocols, or number of members. Its limits are indefinable except for any instance in time since it is in constant flux. It is an open network, available to any that chose to participate. An unbounded network can consist of both bounded and unbounded systems that are subsets of the total network [RE99]. This concept is relevant to NMCI since it will participate with the Internet, a definitively unbounded system, for a relevant portion of its functions and services.

A bounded system, by contrast, is one whose elements are controlled by one central authority that possesses the right to impose policy, sanctions, and can be completely enumerated and controlled. NMCI will be in part a bounded system that is regulated by a central authority, regionally if not globally. Theoretically a bounded system's behavior can be understood by examining its individual parts. Table 1 summarizes the characteristics of the two different views of networks.

Bounded Systems	Unbounded Systems
<ul style="list-style-type: none"> • Centralized administrative control • Total visibility of network nodes • Behavior predictable by examining the components of network. 	<ul style="list-style-type: none"> • Multiple administrative domains with no central authority. • No global visibility (full enumeration not possible) • Interoperability determined by convention • Widely distributed interoperable networks • Users and attackers can be peers in the ecosystem • Can Not be partitioned into finite number of bounded systems

Table 1 Bounded vs. Unbounded Systems

With these definitions we can turn our attention to understanding the concept of Network Survivability. The domain of survivable networks is one that is dominated by large unbounded networks that coexist and collaborate to create a common ecosystem in which they collectively exist and materially participate. The Internet or World Wide Web is exactly this type of ecosystem. Thousands of networks interconnect around the globe to create this mechanism for commerce and information exchange. NMCI will be one of those networks. While bounded in a sense, many nodes within NMCI will have access to the Internet - and vice versa, creating a connection to this unbounded system. In addition, many members of the NMCI community, particularly the surface ships of the Navy, will enter and exit NMCI while participating in other unbounded networks in the interim. The net effect is to give NMCI some characteristics of an unbounded network, making the application of network survivability to NMCI relevant.

What sets survivable network systems apart from simply secure networks is the focus on the central mission of the network. The distinguishing characteristic of a survivable network is its capability to provide essential services in the face of attacks [NRM00]. Identification of the network mission, and the essential services needed to accomplish that mission are then critical to the concept of survivability. Essential services are defined as the system functions that must be maintained to assure the networks mission success, when the system is under attack, suffers failures, or experiences or detects threats. There may be several essential services or several sets of essential services that can be complimentary or duplicative in function. These can be grouped or layered so as to meet the requirement of an essential service through a multiple of methods. It should be remembered that the primary goal of survivability is the accomplishment of the networks assigned mission, not the preservation of any one component, node, or subnet within the network. To maintain the essential services, survivable networks must demonstrate the key properties listed in Table 1-3. These properties are the fundamental building blocks for developing a survivable network system and are the categories that define the survivability services requirements of the network.

Key Property	Description	Strategies
Resistance to attacks	Strategies for repelling attacks	Authentication Access Control Encryption Message filtering Survivability wrappers System diversification Functional isolation
Recognition of attacks and damage	Strategies for detecting attacks and evaluating damage	Intrusion detection Integrity checking
Recovery of essential and full services after attack	Strategies for limiting damage, restoring compromised information or functionality, maintaining or restoring essential services within mission time constraints, restoring full services	Redundant components Data Replication System backup and restoration Contingency planning
Adaptation and evolution to reduce effectiveness of future attacks	Strategies for improving system survivability based on knowledge gained from intrusions.	New intrusion recognition patterns

Table 2 Survivability Characteristics

In the development of the survivability requirements, each category must be subdivided based on the standard attack profile. Typical intruder profiles can be subdivided into three separate phases. They are *penetration*, *exploration*, and *exploitation*. In the penetration phase an intruder attempts to enumerate, profile, and then enter a network through the exploitation of known system vulnerabilities. Once the intruder penetrates the network he enters the exploration phase. In the exploration phase the intruder attempts to further enumerate the network and examine its internal structure for weaknesses. Having successfully penetrated and explored, the intruder has the desired access to the system and begins compromising actions or damage to the network capabilities. Requirements definitions for resistance,

recognition, recovery, and adaptation services assist development of survivability strategies to deal with each phase of an intrusion [RJE99].

Resistance is the ability of a network to deter or repel intruder attempts to penetrate and explore its system. Resistance embodies the majority of traditional computer security. Firewalls, encryption, user authentication, and file access controls are the state of the art for computer security and are the first line in resistance strategies. Diversity is also a resistance strategy and is intended to produce a non-homogenous target set within the network. Creating diversity of operating systems, programs, or network routing mitigates or eliminates the intruder's ability to compromise additional hosts based on a common configurations of identical software. Diversity requirements can be more difficult to achieve because the concept runs contrary to the common business model that demands the economic gains achieved through 100 percent commonality within a system.

Recognition is the ability of the network to perceive and react to patterned or atypical activity that precedes a penetration, exploration, or exploitation event. Recognition strategies are the use of Intrusion Detection Systems (IDS), log parsing, and the use of intelligent agents. IDS typically rely on known patterns of intruder behavior, or through anomaly detection based on the user profile. Intelligent agents work within a host computer and monitor registries for changes in configuration, reporting any changes to a central administrator, monitoring software, or to a "black box" internal to the host for post event reconstruction. Recognition is relevant and critical to all three phases of attack as

recognition at any point in the intrusion is essential for the system to perform any recovery or adaptation services.

Recovery is the ability of the network to reconstitute or restore essential services either during or after an intrusion has occurred. Recovery requirements are what distinguish survivable networks from systems that are only secure [RJE99]. If an attack can not be repelled, a system must have a capability to recover in order for it to be survivable. Recovery is most relevant during the exploration and exploitation phases. Typical recovery strategies are off site data backup and storage, backup or redundant hardware (RAID), host mirroring, and transaction roll back processes. Recovery must also consider the ongoing operation of the network and the maintenance of essential services. The ability to segregate traffic based on the condition within the network and the priority of the individual message is key to maintaining mission essential services.

Adaptation is the ability of the network to rapidly update itself to eliminate exploitation of the network due to poor administrative control. Adaptation strategies include the auto updating features of some software or the updating of intrusion detection rule set based upon published alerts or updates. The limited actions of some firewall software are also an adaptive system behavior. Adaptation requirements relate to all categories of survivability services, as adaptive behavior must be present in each for them to remain effective and relevant in the providing support to the survivable network system. A lack of adaptability would reduce the overall survivability of any individual survivability service as well as the survivability of the network as a whole.

These four concepts are the fundamental elements of survivable networks. They are the yardstick for measuring the capability of a networked system to survive and continue its assigned mission. To apply these principles to NMCI we must first have a basic understanding of the security architecture of NMCI and how the system intends to defend itself.

D. NMCI SECURITY ARCHITECTURE

NMCI is an expansive terrestrial network that is comprised of five essential components. The components to the NMCI infrastructure are a dedicated wide area network, six regional network operating centers, many local area networks, server farms, and client computers or "seats" [RAY01]. A multiplicity of technical protections and policies are deployed in a layered manner to achieve the desired level of information assurance within NMCI. This process is used to achieve a high resistance to attack and minimize the weaknesses of any single security component of the defensive mechanism [RAY01]. There are five basic elements that constitute the NMCI information assurance architecture as defined by the NMCI Information Strike Force. They are [RAY01];

- Network Boundaries and Infrastructure
- Public Key Interface and Directory System
- Seat
- Server
- Security Operations Center

The Network Boundaries and Infrastructure and the Seat are the focus of this work because they embody the majority of the issues related to a survivable network systems.

Changes at this level can improve the survivability of the network without adversely changing the system architecture or changing significant components. If effective survivability measures can be employed within the infrastructure and client seats, then these other architectural elements can be both more effective and more survivable. The Public Key Infrastructure, Servers, and the Security Operations Center will be described in basic detail for understanding of their operation or intent. The Public Key Infrastructure, the Security Operations Center, and Servers are relevant; however, a detailed examination of them extends beyond the scope of this thesis.

The NMCI Network Boundaries are a standardized set of policies and protective mechanisms that define both the interface between NMCI and other networks or an enclave of security within NMCI. These other networks include the internet routing protocol network (NIPRnet), secret internet routing protocol network (SIPRnet), IT-21 networks, Marine Corps enterprise network (MCEN), DISA and commercial WANs, and the Internet [RAY01]. NMCI will interface with each of these and depending on the level of trust deemed appropriate for the collaborating network. A variation on a standard suite of hardware and software are used to achieve information assurance. The boundaries created within NMCI effectively create a series of enclaves that are intended to protect the network and the data that flows within it. The individual boundaries are identified as the Transport Boundary and Boundaries 1 through 4 and each Boundary, with the exception of Boundary 4, possesses both a classified and unclassified side of the network. Each boundary has specific tasks it is designed to perform and the configuration of the hardware systems and the

associated policies reflect the level of trust associated with the collaborating network that is connected at that boundary.

The Transport Boundary is meant to provide protection between NMCI and the wide area network transport services provided by either DISA or a commercial very high performance backbone network services (vBNS). In addition, the Transport Boundary provides the connection for remote dial in services via the UUnet. Protection of the physical assets that make up these two wide area transport services are the responsibility of either DISA or the commercial provider as appropriate. The vBNS and DISA incorporates the following protective features **[RAY01]**:

- Denial of Service Protection
- User Data Confidentiality
- Identification and Authentication
- Access Controls
- Security Alarms and Audit Trails
- System and Data Integrity
- Personnel Security
- Physical Security
- Ongoing Security Improvements

As with all the other boundaries within NMCI, the Transport Boundary possesses both a classified and unclassified side. The unclassified side of the wide area network relies upon virtual private network (VPN) devices, IP layer protections, intrusion detection systems, and policy based routing for protection.

The protection mechanisms of the Transport Boundary are positioned at the access points of the WAN. Virtual Private Network (VPN) devices, routing table

authentication, and IDS monitoring compose the defensive elements of the unclassified network. For the secure portion of the network, Type 1 encryption is used and the IDS are omitted. Host IDS systems guided by network security policy are deemed adequate to prevent illegal activities at the host level (Boundary 4) and the bulk encryption used will prohibit access by other vBNS and DISA users unauthorized access to NMCI [RAY01].

Boundary 1 provides protection between NMCI and any external network, to include NIPRNet, SIPRNet, and the Internet. The mechanisms used for defense are firewalls, content scanners, IDS, routing table authentication, and both single and dual sided VPN service. Classified use of boundary 1 is via the (SIPRNet) and Type 1 encryption is used for communication across the SIPRNet. Single and dual sided VPN is also available through Boundary 1 in a manner consistent with that of the unclassified Boundary 1. The primary difference between the classified and unclassified side of Boundary 1 is that only one firewall and one content scanner will be used on classified side. This is deemed adequate since through-put for the classified side is anticipated to be substantially lower [RAY01].

Boundary 2 provides protection between legacy systems and NMCI. The definition of a legacy system is applicable to BAN's and LAN's that were deployed prior to the inception of NMCI and its security policies. This specific definition of legacy systems includes all IT-21 networks (shipboard), the Marine Corps Enterprise Network (MCEN) and other legacy base area networks and application that reside within the Navy organization or are accessed by a part of the Navy organization for operation. The defensive mechanisms here are the same as used in Boundary

1. Boundary 2 is also has both a classified and unclassified side to it. By default access to Boundary 2 will be via the firewall suite. However, it is anticipated that some systems will not meet firewall policy and therefore be routed via a VPN connection. In this situation the legacy server would remain within the legacy network and the NMCI user would be connected via a VPN client. The access to the legacy applications carries with it some risk that should be balanced with the functionality gained by their entrance into NMCI. These legacy applications often contain know vulnerabilities and if compromised could provide a point of entry for an adversary into the NMCI environment. The Boundary 2 requirements for legacy applications are still to be determined and ultimately will be approved by the NMCI Connection Approval Process [RAY01]. Classified Boundary 2 design would be similar to the unclassified design [RAY01].

Boundary 3 provides protection between communities of interest (COI's). How NMCI deals with communities of interest is based on their sensitivity and the geographic location of its members. The defensive mechanisms used for each COI is then based on this same information. Table 1- 3 summarizes COI's and what mechanisms would be used.

	Group	Virtual LAN (VLAN)	Shared VPN Gateway	Dedicated VPN Gateway	IDS	Firewall
Sensitive (A)	X	X ¹				
Highly Sensitive Distributed (B)	X	X	X	X ²		
Highly Sensitive Co-Located (C)	X	X			X	X ³
Isolated (D)				X	X	

Notes: 1. To limit network access to a private server
2. To protect private server or enclave with its own LAN
3. If required

Table 3 Communities of Interest within NMCI

For COI type A, enforced group policies and potentially a VLAN are used to control access to their server. For COI type B, VLAN, policy based routing and would segregate the COI from NMCI. VPN's and group policy would provide data confidentiality and access to distributed members of the group. For COI type C, the server resides in the same location as the members. Group policy and a VLAN control access to the server and a firewall may be deployed to improve the segregation of the group from NMCI. For COI type D, NMCI provides only the connectivity to the server for data access. In this situation a dedicated firewall and an IDS are deployed. NMCI will also support foreign nationals who are assigned to the Department of the Navy (DoN) installations, activities, or commands within Boundary 3. Specific configurations are made for their use within NMCI [RAY01].

Boundary 4 protects NMCI at the host and server level. The defensive systems employed at this level are numerous and are dependent upon the classification of the host or server. They include but are not limited to secure

operating systems, VPN client, Smart Card sign-on, email encryption, web server authentication, host IDS, virus scanning, and policy enforcement [RAY01]. The specific configuration of the host or server can be determined by examining the CLIN for the specific seat being used. The CLIN configurations define exactly what classification the seat is cleared for and what the configuration of the seat should be.

The public key infrastructure for NMCI will be developed and managed by the Department of Defense (DoD), with the National Security Agency (NSA) and DISA responsible for the development of the core components. The DoD will establish a central certification authority to create, assign, and issue public key certificates for NMCI. The same organization will maintain the directory. The directory will be based on Windows 2000, using the Windows 2000 Blackcomb update, identify and authenticate for the domain or logon.

NMCI security is managed and controlled through NMCI security operations centers (SOC) that are co-located with the regional NMCI Network Operating Centers (NOC). The six classified and six unclassified SOC's monitor the IDS, manage firewall policy, virus and content scanning, encryptors, VPN devices, and remote access servers [RAY01].

Client Seats and their configurations are numerous, and the possible number of variations significant. The Contract Line Item Number (CLIN) list contains all CLIN information and the specific descriptions of each client seat. The primary element to be taken from the seat configuration is the importance of standardization of each seat, server, router throughout the NMCI network. All hosts will run a common Windows-based operation system and

a common software suite. Windows will also be the basis for the authentication and identification for server operation. The last major component of the security architecture to be examined is the Network Management Network (NMN).

The Network Management Network is the mechanism for the monitoring, updating, and configuring the routers, servers, and switches that reside on the unclassified side of NMCI. The NMN transport is provide through a separate wide area network that connects the NOC's with the individual routers, servers, and switches. The status of these nodes will be monitored via the NMN using HP Openview, Tivoli, and Remedy software suites. The NMN has no redundancy or fail-over and therefore when connectivity is lost, monitoring service capabilities will also be lost until service is restored [RAY01]. Having completed a lengthy look at the major components of NMCI Security Architecture, we can step back and begin to examine the potential weaknesses that arise from it.

In examining the survivability of the NMCI Architecture we should begin by parsing out the overall network into three major areas of interest. Those three areas are availability of the network (A_o), security, and quality of service (QoS) (or differential service). For a system to be survivable it must be available, secure, and possess differential services that permit traffic to be segregated or prioritized when the network is under stress or in extremis. Taking a look at the NMCI architecture in this manner will help understand the difference between a secure network and a network that is both secure and survivable.

E. AVAILABILITY

To evaluate the availability of a network we must first define what availability means and then consider the basic principles of high availability engineering. When we have done these things we can then use them as a framework to compare to the NMCI structure. Availability of a network has two different definitions that can be used to evaluate performance. In a telephone circuit, network availability is defined as the ratio of the time the circuit is operational to total elapsed time. In a network switching system availability is defined as the accessibility of input and output ports. For the purpose of this thesis, availability will use the later definition, that is the availability of access to input and output ports. This definition is relevant since NMCI will be essentially a stateless system. NMCI is not concerned about the state of a particular connection, but intends to provide connection end-to-end utilizing a packet switched network. Since state is not a consideration, the important aspect of NMCI availability is then the operational capability of the hardware systems that make up the NMCI infrastructure. Under the NMCI contract the Navy and Marine Corps will not own the hardware infrastructure. The Navy has contracted for services from a vendor and the desired availability for services was agreed upon in a conforming contract awarded in October 2000. The best way then to evaluate the A_0 of NMCI is to compare the contractual agreements with the basic principles and practices used in the construction of networks. The key elements to high reliability engineering are to eliminate single points of failure, provide reliable crossover (from primary to backup), and promptly detect failures upon

occurrence [RB02]. Appendix 1 contains a discussion and examination of how a multi-threaded system can achieve availability rates of 99.99% or above. This can be achieved through simple redundancy of components within the supporting infrastructure.

Independently generated statistics on the availability of commercial network providers can be found on the internet. These statistics show that commercial only 6 of the 26 providers surveyed by this site could not meet 99.9 % availability (reachability) [MTX02]. The average among all the US providers from this survey was 99.86%. When compared to the standard performance measure contained in the NMCI contract, 99.86 % availability met or exceeded the requirements for each specific service level agreement [PEOIT02]. The NMCI contract also includes provision for evaluation of services based on latency and packet loss within some of the SLA's. Examining the same open source data we can see the average latency was 67.12 mili-seconds and packet loss was .2705%. When compared to the standard performance measure contained in the NMCI contract, .2705% packet loss and a latency of 67.12 mili-seconds exceeds the requirements of the relevant service level agreements [PEOIT02].

In essence, the Navy has contracted for services that are no more reliable than what is presently available from any other internet service provider (ISP).

While this may be adequate for some services used within NMCI, it may be inadequate in times of crisis, or for high integrity or time critical data used in battle management or force projection operations. Any limitations that result from this are then relevant to what NMCI should be used for and ultimately what the mission function or

functions of NMCI should be. They also go directly to the overall survivability of the network. Even if security and quality of service are established within a network, the system availability must be adequate to meet the needs of the primary mission. The SLA's negotiated for the service should be traceable to the operational requirements for the network. If the network availability is inadequate to support the mission essential functions then the network has failed.

F. QUALITY OF SERVICE

Equally important to mission functions of NMCI is the concept of quality of service. Again we must begin with a definition. Quality of service refers to the ability of a network to provide better service to selected traffic that is flowing within the same network. The primary goal of QoS is to provide priority including dedicated bandwidth, controlled jitter and latency (required by some real-time and interactive traffic), and improved loss characteristics. Also important is making sure that providing priority for one or more flows does not make other flows fail [CSCO02]. This can be contrasted with the current methodology within the internet of "best effort" delivery. Under best effort delivery methodology, all packets (or transmissions) are treated equally and delivery to their destination is arbitrated by routers on the pathway the packet takes to its destination on that equivalent basis. When the network experiences congestion due to demand, failure, compromise or a combination thereof, all users experience an approximately equivalent level of degradation. This is suitable for commercial and most military applications, but is inadequate for military

applications that demand as near real-time or "hard real-time" service. Some tightly coupled combat system applications require a deterministic level of service that guarantees delivery within specific time frames to an end user [RB02]. Deterministic service then includes the idea that the packet has guaranteed delivery to an end user within a specified time. To achieve this within a network we then must consider that determinism has two levels of complexity within itself. First, we must establish determinism within the application. How do we assign the priority to packets of a specific application? Second, we need to resolve the priority among a multiple of "priority" applications flowing within the network [RB02]. Policy and configuration solutions, or partial solutions to these problems exist, but they are not widely deployed in either the commercial or military world because of the complexity and effort required to implement and manage them. At present, the standard practice to resolve quality of service issues is the massive application of bandwidth because of its relative low cost and ease of implementation.

While the deployment of more bandwidth is a solution, it is a hardware solution that in a time of crisis may not be adequate to a combined or coordinated attack on a network. The topology of a particular area may limit the alternative paths available. A series of cable paths buried in a common trench or terminating at a common point are single points of failure within the network. For a network to be survivable it needs to be able to discriminate among the traffic flowing within the network and handle that which is mission essential while delaying or discarding that which is not. The inability to perform

this discriminatory behavior results in default "best effort" service to all users, granting equal privilege to traffic that is both mission and non-mission essential. For this reason there should be another solution to the problem of providing a deterministic level of service within NMCI for it to be a survivable system.

G. SECURITY

The term security contains within it a long list of characteristics that express what designers hope to achieve within any particular network. The presence of these characteristics determines the level security that exists within that network. Consequently, the lack of their presence also says a great deal about how secure the network should be considered. The characteristics that define security within a network are;

- Confidentiality. Unintended recipients can't read our traffic. Confidentiality includes secrecy of the data.
- Authenticity. Unintended originators can't fake traffic. Nobody forged my messages. Authenticity is a superset of integrity.
- Integrity. Traffic hasn't been tampered with. What you got is what I really sent you.
- Non-repudiation. I can't get away with saying something and later denying it.
- Access control. Unauthorized users can't use network and computing resources. More colloquially, keep the riff-raff out of my corner of the 'net.
- Assurance of service. The network is available for use when I need it. Resistance to denial of service attacks.

- Traffic analysis. Ability to derive intelligence from the addresses of messages, even if the contents are confidentiality-protected.
- Traffic flow analysis. Derivation of intelligence inferences by observing flows to and from commands and individuals.
- Interceptability. Ability of unintended recipients to receive traffic (regardless of whether they can read it).
- Jammability. Vulnerability of a link to interruption by signal interference. **[RB02]**

There are numerous ways of achieving these characteristics within a network and they can be applied to a multiple of layers within the OSI model. This is drawn out in ISO 7498-2 that lists the potential areas of application of security measures at each of the seven layers within the OSI model.

Service	OSI Layer
Confidentiality	1, 2, 3, 4, 7
Authentication	3, 4, 7
Integrity	3, 4, 7
Access Control	3, 4, 7
Non-Repudiation	7

Table 4 ISO 7498-2 Layer Model

It should be noted that the ISO has expressed these as theoretically possible at the noted layers of the OSI model. ISO 7498-2 makes no reference to whether the decision to choose any one or all of the relevant layers is either effective or manageable. It should also be noted that there are specific characteristics (traffic analysis,

traffic flow analysis, interceptability, and jammability are not even addressed). Taken a step further we can look at how the OSI model overlays with the requirements and solutions in what I shall refer to as the "Buddenberg Matrix". This table better illustrates the objectives, methods, and examples of how the security characteristics can be dealt with within a network [RB02].

ISO RM Layer	Requirements	Solution	Objective	Examples
7	Confidentiality Authenticity	Object Level Security	Object Level Security	S/Mime, secure shell, secure socket layer , VPN
3,4	Perimeter Protection of Enclave (Prevents DDoS attack)	Firewalls Intrusion Detection MAC/DAC	Secure the Network Box (not the data)	Passwords Firewalls IDS
1,2	Traffic Analysis Traffic Flow Analysis Jammability Detectability	Link Crypto LDI/LPD Spread Spectrum	Secure the Network pipe(transport)	KG-84 STU-III Wireless LAN

Table 5 Buddenberg Matrix of Security Requirements

The operative theory behind the Buddenberg Matrix is that in order to achieve the most efficient and effective level of security, all the elements (problem, solution, objective, and application) must be in alignment. This is not to say that security measures are not or cannot be employed in another fashion. What this matrix demonstrates, however, is two things;

(1) misalignment at best creates significant inefficiencies in network implementation and operation, and

(2) misalignment at worst creates security which misses the objective and in actuality provides less security than desired.

Using the Buddenberg Matrix and the logic derived there, we can examine the security architecture of NMCI to evaluate the logical match between the requirements and solutions needed to secure NMCI.

Recalling the earlier discussion and description of the NMCI security architecture we readily see that enclave security is the centerpiece of its network defense scheme. The defense-in-depth strategy employed by NMCI is implemented in physical and logical layers to provide security enclaves at the regional, command, and host level. It is intended to provide confidentiality, integrity, availability, non-repudiation, access control, authenticity, identification, survivability of information systems [RAY01]. Immediately there is an obvious mismatch between the requirements and the methods employed when the NMCI architecture is compared to the Buddenberg Matrix. The defensive mechanisms are employed primarily at layers 3 and 4 of the OSI layer model while some of the requirements can best be handled at layer 7, the application layer. This indicates that more efficient and effective security may be had by employing object level security at the application level. When compared to the ISO 7498-2 recommendations in table 1-4 we can see that many of the security characteristics can be addressed at OSI layers 3 and 4, but non-repudiation should be handled at layer 7. The security solutions employed within NMCI may not provide the most effective security because of a mismatch between the requirements and the applied technologies. Looking at confidentiality and authenticity we can see the mismatch by referring to the Buddenberg Matrix.

ISO RM Layer	Requirements	Solution	Objective	Application/Methodology/Technology
7	Confidentiality Authenticity	Object Level Security	Secure the data	S/Mime, secure shell, secure socket layer , VPN
3,4	Perimeter Protection of Enclave (Prevents DDoS attack)	Firewalls Intrusion Detection MAC/DAC	Secure the Network Box (not the data)	Passwords
1,2	Traffic Analysis Traffic Flow Analysis Jammability Detectability	Link Crypto LDI/LPD Spread Spectrum	Secure the Network pipe (transport)	KG-84 STU-III Wireless LAN

Table 6 NMCI Misalignments within Buddenberg Matrix

Within NMCI, confidentiality and authenticity are being handled primarily by the application of various encryption methods. This constitutes a misalignment between the requirements and the technology when compared to professor Buddenberg's matrix. VPN's do reside within NMCI and could provide the authenticity and confidentiality desired, but their application is less than optimal.

As designed, NMCI uses VPN's to provide access to legacy applications and servers, and for remote access [RAY01]. This application provides access to these systems as an alternative to the weakening of firewalls. This application, however, is limited in scope and doesn't provide end-to-end security. In the context of this thesis, and in reference to the Buddenberg Matrix, end-to-end security means that the IP datagrams travel from the sender to the intended receiver host in a "black" or encrypted form. The originating and destination IP addresses may or may not be encrypted, but the data within the message is encrypted by the sender and remains so until decrypted by the receiver.

The figure below is taken from NMCI System Security Authorization Agreement and depicts the Classified Boundary 1 configuration. Note the VPN path is depicted in the dashed lines connecting the VPN clients external to the network to the NMCI classified BAN.

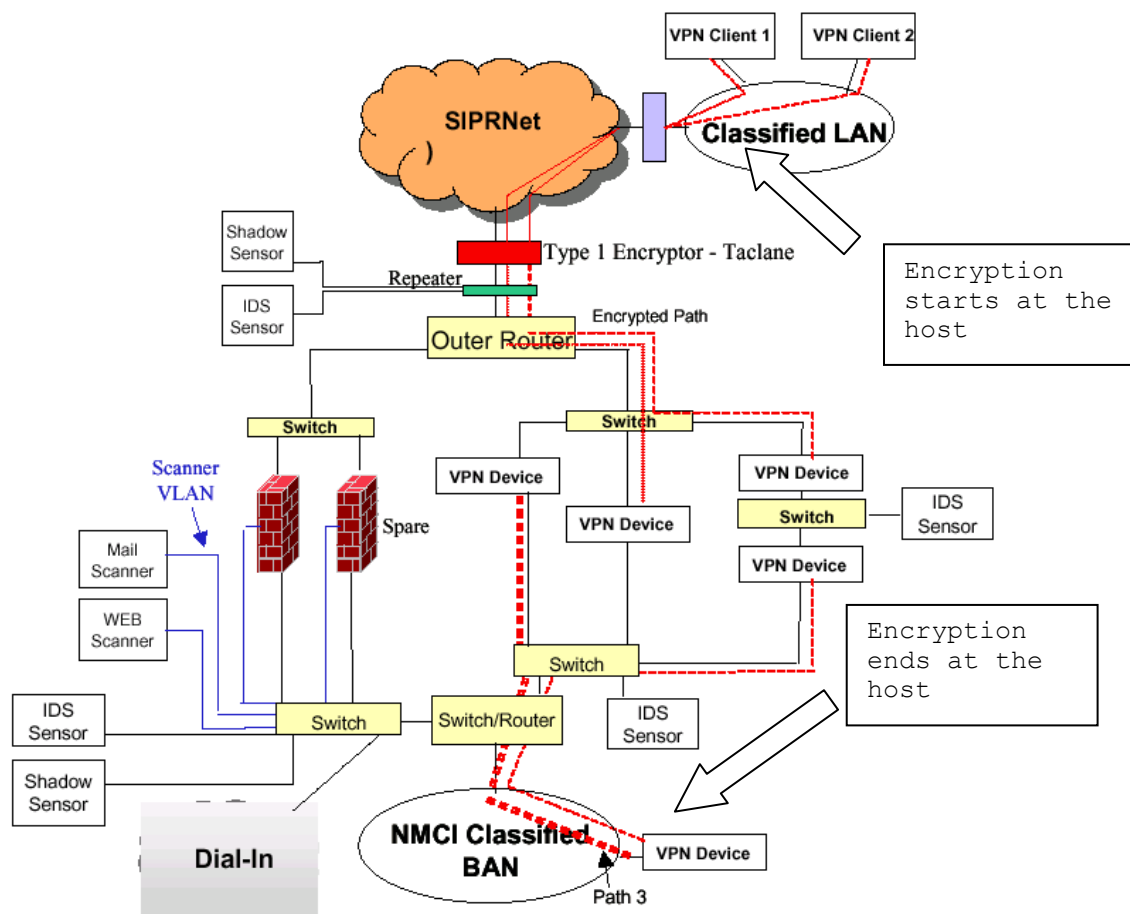


Figure 2 Classified Boundary 1

(Figure 2.1.2-2 from NMCI System Security Authorization Agreement of 19 March 2001)

The limitation to this application is that the encryption is provided only as far as the gateway. Once the VPN reaches the gateway the information travels in an unencrypted form within the network and is stored within the network servers

in the same way. The same configuration is used within other Boundaries within NMCI. Below is a diagram of the unclassified Boundary 2. The encryption ends at the VPN device. The weakness of this is that there is no consideration given to the potential alteration of the message from within the NMCI network. So while the authentication and confidentiality can be maintained between the client and the gateway, there is no such guarantee between the gateway and the NMCI host that that reside within the BAN or LAN

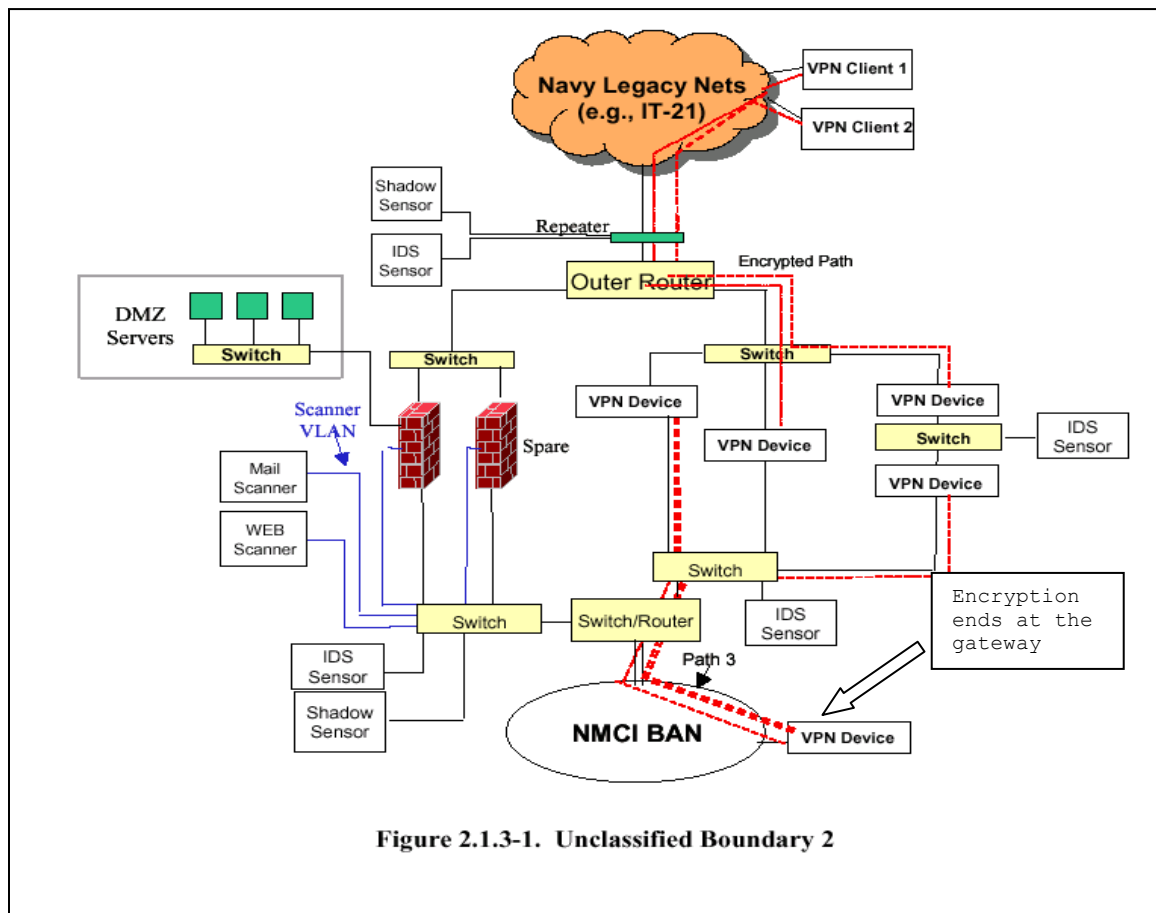


Figure 2.1.3-1. Unclassified Boundary 2

Figure 3 Unclassified Boundary

(From NMCI System Security Authorization Agreement of 19 March 2001)

The assumption is that the data would be inaccessible or that an intruder would not be able to enter the NMCI BAN

without being discovered. Insider attacks are a real threat to any network and this application of a VPN, while providing some security, also presents a vulnerability in its application within NMCI. It should be noted, however, that the application of VPN's, as well as many other devices to provide network security, are merely the tactics employed in the battle for network security. The network defensive strategy, or the underlying logic for the use of the specific tactics within the network, must also be examined to obtain a full picture.

As discussed earlier in this chapter the NMCI network defensive strategy is essentially an enclave-based system that relies upon a multiple of layers to restrict unauthorized access or malicious behavior within the network [RAY01]. This enclaving strategy represents the state of the art for network defense as it is employed today. The problem with this strategy is that historical data has shown us that making a network a hard target alone will not ensure its success. The CERT data listed in Table 1-1 on page 7 of this text shows that for all the effort made, security compromises still occur. The strategy of enclaving networks assumes that security is a binary relationship. The multiple layers are intended to blunt or repel the attack, and the system will not be compromised. If, however, the attack succeeds, the enclaving strategy does not address what to do next. Under this strategy networks are either secure or compromised, and there is really no middle ground. There are usually plans for contingencies, they tend to revolve around the recovery and reconstitution of the data within the network in a post event environment. This is a non-real time evolution that does not provide for the continued operation of the

network. NMCI is considered a mission critical system and its loss would have a serious effect on the mission support and operation of the Navy and Marine Corps [RAY01]. To fail to provide continued operation cedes victory to those that hope to conduct "information denial" attacks against NMCI.

The concept of information denial was drawn from the theories of Admiral Alfred Thayer Mahan and follows his theories on sea control and sea denial in times of conflict. Information denial is the practice of disruptive or destructive activities intended to deny the flow of information across the network [RB02]. The enclave strategy only partially addressed information denial attacks by providing the best possible resistance within the network. Once penetrated, however, the defensive enclave provides little or no capability to ensure continued operation. This is particularly true when the homogenous software suites intended for NMCI are considered. The standardization of software suites across the network both improves security while providing a common weakest link. A compromise employed at one nearly guarantees that exploit will be successful at any other similarly configured point within the network.

In Mahanian terms, the NMCI enclave strategy hopes to ensure information flow across the entirety of the network by securing regional "information dominance" through the defeat of information denial attacks. Much like the strategic hamlet strategy of the Vietnam War, the initiative is ceded to the attacker, allowing him to pick the time, place and method of attack. In network terms, the attacker need only find a single crack to enter, while the defender must protect all avenues of approach, even those he would be

unable to see. Applying the Survivable Networks Systems strategy would be employing Mahan's concept of sea control to the network environment, "information control" [RB02].

The strategy of information control is the process of transmitting traffic across the network in the face of all the information denial activities. Implicit to this is the guarantee of authenticity and confidentiality, ensuring the integrity of router configurations, and planning the network to be appropriately robust to guard against individual link outages [RB02]. Recall that Network Survivability is defined as the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents [RE99]. The concepts of information control and Survivable Network Systems are therefore closely related, if not identical in intent and purpose.

Just as the U.S. Navy practices sea control as its basic strategy, NMCI should employ information control to its network operations. Employing such a strategy aligns the tactics of network defense with the network strategy of continuous network operations. Designing the system to operate in spite of damage or compromise is essential to the mission success of NMCI and the U.S. Navy, and therefore the application of the Survivable Network Systems concept is relevant to security of NMCI.

H. METHODOLOGY FOR EXAMINING NMCI

The examination of NMCI conducted in this thesis is in every sense on a "macro" level. The Survivable Networks Analysis method was used to guide the examination. The focus of this thesis is to give the readers a larger view of the NMCI structure as a whole. Many individuals inside

and outside the Navy organization are working diligently to find solutions to extremely complex and complicated problems within NMCI. There appears, however, to be a lack of focus on the larger context of what NMCI will mean to the Navy and how this larger view of NMCI relates to the very specific problems being addressed.

The Survivable Networks Analysis method was applied to NMCI in the following manner. The first step was to gather information available from DoD and business sources on the structure, composition and mission of the Navy and Marine Corps Intranet. All the sources used for this part of the analysis were publicly available from either the primary contractor for NMCI (EDS Corporation) or the relevant DoD web sites. These materials formed the basis for evaluating the mission and business model of NMCI under the Survivable Networks Analysis method.

Next I gathered information from the end users of NMCI about what they desired to achieve from the implementation. These client organizations provided input as to what they were presently engaged in over networked systems and what they would like to see provided from NMCI. These interviews were used to help evaluate the mission of NMCI and develop what the essential services might be under the Survivable Networks Analysis method.

Third, a top level view of the NMCI security architecture was made and the system evaluated for soft spots, vulnerabilities, or a lack of capability when evaluated under the Survivable Networks Analysis framework. Fourth and last, an overall evaluation of the survivability of NMCI was made, based upon the information derived from the previous three steps. When summed up, the total effort should provide a top level view of the NMCI system and what

its mission, essential services, vulnerabilities, and overall survivability are under the survivable network analysis method.

II. MISSION AND THE NAVY AND MARINE CORPS INTRANET

A. THE CURRENT MISSION DEFINITION

The network survivability analysis method first examines what the mission of NMCI is and from there determines what the mission essential services for NMCI are. By defining the mission function for NMCI we can then determine what essential functions within the network must be maintained for the system to remain viable and function in the desired manner in a time of compromise or when under attack. As discussed in Chapter One, the network is intended to provide connectivity throughout the entire force in cooperation with, and as part of, the Global Information Grid [RD00]. NMCI is essentially a communications system that connects all the DoN activities in a manner similar to other communications systems that already exist. The obvious differences between NMCI and the other communications systems are in the mission definition and scope of the communication that occurs within them. NMCI, unlike the other communication systems, has as its mission the maintenance of connectivity of all the participating members.

The majority of the communications systems or networks used within the Navy and Marine Corps operate with a purpose or mission function as their focus. Communication plans within a battle group, amphibious group, or expeditionary force are designed around specific functions or mission areas with the intent of segregating voice traffic and data traffic into logical subsets for ease of use by the operators. When decision makers request information, or pass it on to others, there is typically a

dedicated circuit that correlates to the mission area relevant to the information. Consider the following example of a communication network that is part of a communication plan with a battle group.

1. Alpha-Xray Battle Group Communications Example

The responsibility for the coordination of anti-submarine and anti-mine warfare resides with the undersea warfare commander (designated AX) within the battle group. His job is to coordinate the employment of anti-submarine and anti-mine warfare assets (ships, submarines, and aircraft) for the battle group commander in both offensive and defensive operations. To complete his mission AX will have two or more dedicated radio circuits available to employ his undersea warfare assets. These radio circuits are segregated by function and AX is responsible for arbitrating the use of these by his warfare elements based on his (AX) prioritization scheme. For example, one circuit can be used for contact and reporting, and another for coordination. If there are adequate radio circuits available, these can be subdivided by region or relative position from the battle group, further segregating the traffic (subnets). Within each of these subnets, the reporting information is prioritized in a predetermined manner by the warfare commander. Actual contact reporting has the highest priority, while other forms of communication fall in an established hierarchy based upon the status of the network and the level of activity.

During periods of high demand or limited connectivity, some voice calls are not required or not permitted at all. The reporting and coordination information is often further

segregated in an ad-hoc fashion by the participating units, who may randomly select an unassigned or low demand frequency or channel for coordinating communications. An example of this is the use of 303.0 MHz frequency, commonly referred to as "Winchester", for local area coordination by the assigned assets. A means of fail-over is provided to the participating assets through a predetermined hierarchy among the channeled frequencies.

If a channel fails or is unusable, the undersea warfare units have a predetermined routing for their information by merely switching to another dedicated circuit for passing priority information. If the undersea warfare circuits fail completely, the members do have the ability to use other nets that are dedicated to other mission areas if their traffic is deemed to be of a high enough priority by the warfare commanders that control those other mission area subnets. There are often other communication systems, some intended primarily for data, which contain a voice channel that can be used for voice communication in local area coordination. The LAMPS Mark I VHF data link and the LAMPS Mark III microwave tactical data link are examples of data systems that can be used for voice communication with other assets.

Regardless of the channel or channels utilized by a reporting unit, there is an accepted standardized format for the establishment of communications and the content of each transmission within any subnet. Each transmission begins with a call to the destination, giving the destination call sign and followed by the sending unit's call sign. The receiving unit then acknowledges that call in the same format, giving the sending unit permission to begin their data transmission. The format of the

transmission and the method of establishing communications can be viewed as a very crude use of internet protocol datagram employing the three way handshake to establish communication.

When viewed in the aggregate, the well established communication plans and protocols used by the Navy and Marine Corps for radio frequency communication and the standards for IP networks have many similarities. Each node of the network has a distinct address (call sign) used for communications. Both networks have a protocol for the establishment of communications between two or more of its members. Both networks have an agreed upon format for the transmission of data. Both networks possess mechanisms or protocols for fail-over protection. Both networks use subnets to segregate communications into relevant communities of interest. For all these similarities there are, however, significant differences that make the IP network less capable when compared to the simple radio frequency nets.

The most significant of the differences is the mission orientation of an IP network, in this case NMCI. The primary mission focus for NMCI is connectivity of the Navy organizations. The implication is that unlike the radio networks there is no established hierarchy or prioritization of communications within the network. There is also no single arbitrator (warfare commander) responsible for establishing a hierarchy or prioritization of communications. The resultant behavior of IP network communications is contingent upon every node getting best effort service. The IP network will attempt to retain connection to all of the nodes regardless of the demand and ignoring the relative value of the information being

transmitted within it. With the focus on connectivity, every transmission has equal value regardless of war fighting or mission relevance. While fail-over protection exists within an IP network, its value is diminished without an accompanying prioritization scheme to assist in the segregation of traffic within the network. While this may not cause a complete system wide collapse of an IP network, it could impair its performance regionally and reduce its usefulness and effectiveness as a communication system.

These differences are significant when considering the numerous NMCI sites located outside the continental United States. NMCI serviced sites in Alaska, Hawaii, Japan, Iceland, Guantanamo Bay, Cuba and Puerto Rico would likely rely on satellite transmission for connectivity to the mainland and the remainder of the network. These sites are geographically isolated and the application of additional optical fiber to increase available bandwidth has its limitations and vulnerabilities. There are physical and fiscal limits to achieving diverse paths on the ocean floor. Radio frequency systems (satellite or terrestrial) do not yet have the bandwidth available to absorb the entirety of the traffic within NMCI for these locations. Without arbitration or prioritization of the traffic within NMCI, these sites could experience significant delays or their respective portion of the network could become isolated by an intrusion or other event. Mission critical communications could be thwarted completely because of limited available bandwidth and the lack of prioritization of traffic. Given that NMCI is considered a mission critical system for the Navy and Marine Corps; this is likely a very undesirable situation [RAY01].

The other major difference between the radio networks and NMCI is the lack of a central authority to arbitrate the priority of traffic. The absence of an arbitrating authority is a direct consequence of the mission focus of connectivity. The decision to make connectivity the focus for NMCI performed the arbitration function in advance, making all transmissions equals of one another. Connectivity is the essence of developing and fostering the growth of community within NMCI and the Navy as an enterprise [JH97]. However, it is probably an inadequate mission focus for NMCI. Even within the internet market forces have begun to push the focus of service beyond that of purely connectivity.

Internet service providers are offering improved quality of service in some areas of application (video and voice) to those users who are willing to pay a premium for it [SBC02]. This is classic market economics at work. But what then does it say about the NMCI mission focus? The implication is that only those commands who can achieve the desired funding will achieve better than best effort service. This is not practical or even logical when considering the established process within the military of developing requirements.

Requirements for military systems of any type derive from the mission need statements generated by those responsible for managing the relevant communities within the Navy or any of the other Services. The communities that exist within the Navy are largely drawn along the lines of warfare specialty or mission area. There is no warfare specialty or community with a single claim to NMCI, and so no community leaders to determine the prioritization for the traffic flow within NMCI. The requirements

methodology with a mission focus should be applied to NMCI given that the network has been identified as a mission critical system by the Navy [RAY01]. What the Navy and Marine Corps team then needs to do is define the warfare mission for NMCI and give it context for its inclusion in the Navy and Marine Corps measure of combat capability.

B. NMCI MISSION DEFINITION

The Navy needs to define the mission or warfare function of NMCI. The "mission need" to connect the Navy nodes is obvious and NMCI will fulfill that need when fully deployed. The Navy needs to take the next step and begin to define just how this massive network fits into our war fighting capability. This process starts by creating a more granular mission definition to NMCI that will allow for the development of a mission hierarchy. Logically then, the arbitration and quality of service issues should fall in place along the lines of the mission capability that NMCI is intended to provide for the Navy and Marine Corps team.

In addition, the responsibility for policy or hardware implementations necessary to perform the arbitrating function will logically fall to the relevant warfare area or support organization commander(s). Ultimately, this method should produce the prioritization of transmissions within NMCI and the application of quality of service that is aligned with both the capability and primary mission functions of NMCI. Redefining NMCI's mission will begin to transform it into the war fighting asset that the Navy's simplistic voice circuits have proven to be for over 50 years. Once this has been done then we can move to define

what the essential services of NMCI are or should be and then how survivable the network is.

C. REDEFINING THE MISSION FOR NMCI

In the process of redefining the mission for NMCI we first need to take a step backwards and begin by looking at what the potential mission areas could be. By defining the potential mission areas they each can be examined individually and their respective requirements mapped to the capabilities of NMCI. This allows for a comparative examination of the each mission individually. The respective mission areas can then be evaluated and the most appropriate one(s) selected through a process of elimination.

To begin the process we must define what I believe are the three potential mission areas are for NMCI. They can be broadly categorized as administration, force projection, and battle management. These mission areas embody the essential functions of the Naval Service and capture the functional requirements necessary to the determination of their applicability within NMCI.

D. EVALUATING NMCI MISSION CAPABILITIES

Given that NMCI as a complete entity does not yet exist, we must make some assumptions so that we can establish a frame of reference for the examination. First, since no historical data exists for NMCI we need to use a surrogate for performing the evaluation. NMCI as proposed today (approximately 360,000 client seats) in size and scope will be comparable to large commercial ISPs [RAY01].

These commercial wide area networks are constructed in a similar manner, using similar technologies and software, and operate in a comparative environment to NMCI. Their performance data then is a valuable yardstick for measuring the future performance of NMCI in each of the mission areas.

Second, a basic assumption about NMCI security must be made since we don't yet have metrics or data for security performance within NMCI. We can assume that NMCI will embody much of what has been gleaned from the over ten years of commercial and military experience in operating wide area networks. The firewalls, IDSs, and virus checkers represent the best of breed in the commercial sector. This assumption is in line with the basic concept of securing a services contract for NMCI. The prime contractor is responsible for security and the assumption is they will provide the best possible security given the monetary incentives provided by the government in the NMCI contract award. To summarize then, the basic assumptions for this comparison are;

(1) Once fully deployed, NMCI will be an equivalent in size and scope to the larger commercial ISPs operating in the United States and that the commercial ISP data is representative of the minimum performance NMCI will achieve when fully implemented.

(2) NMCI will possess equivalent or better than network security compared to that of the larger commercial ISPs operating in the United States and the security capability of commercial ISPs is the minimum performance NMCI will achieve when fully implemented.

In addition to these assumptions we must establish common characteristics for examining performance

requirements of NMCI in each mission area. The relevant factors that determine successful completion of the required tasks in each mission area are analogous to those needed to function in the existing e-commerce environment. In the commercial sector those factors are defined as availability, security, and non-repudiation [RC98]. These factors are essentially identical to those used for the evaluation of the NMCI architecture. For this reason, each mission area will be examined for suitability within NMCI based on the same framework used to examine the NMCI architecture, that of availability, quality of service, and security. This is a logical division of mission requirements and will provide a consistent reference for comparing essentially dissimilar missions.

With these assumptions and common framework established as the ground rules we can turn to publicly available statistics on commercial ISPs and compare them to the requirements of each specific mission area in the process of examining NMCI. If the existing commercial network performance is adequate for the mission requirements it is logical to then assume NMCI performance in the same mission area will be equivalent or better. The converse can also be said to be true. This comparison will form the basis for determining whether a particular mission area is suitable for use within NMCI.

E. ADMINISTRATION

For the purposes of this examination, administration is defined as the functions necessary to complete the day to day operations and maintenance of the Navy and Marine Corps as an enterprise. This includes the coordination and execution of all actions required for the routine movement

of personnel and material, as well as the requirement to transfer monetary instruments to organizations both internally and externally for the conduct of routine, peacetime, non-combat operations. The Administration mission area requirements are analogous to e-commerce requirements in the private sector of the economy.

1. Availability

In a network switching system availability is defined as the accessibility of input and output ports. The requirements for availability in this mission area are essentially identical to that required for commercial ISPs. Commercial networks seek to connect end users for the purpose of communication or commerce. The goal for NMCI under this mission area would be to connect all Navy organizations with their desired end user, whether they are a Navy, other Service, other governmental organization or private company. The NMCI contract award contains specific performance parameters for the availability characteristics of latency and availability. These parameters are established under several separate service level agreements (SLA) that govern different portions of the NMCI services contract and the requirements differ slightly among them. SLA 10 governs NMCI Intranet Performance and is the most germane since it establishes performance requirements for the entire network. SLA 10 therefore will be used as the standard for comparing performance of NMCI to the commercial ISPs. The standard of for availability and latency under performance under SLA 10 are 99.8% and 70-100 milliseconds respectively **[PEOIT00]**.

The source used for statistics on commercial ISPs is by Matrix.Net, an independent web site that monitors many internet providers in the U.S. and around the globe.

Specific data for latency, packet loss, and availability are kept for 26 different ISP's that operate within the United States. Looking at the statistics calculated by this site we can see that 20 of the 26 major U.S. ISP's provide connectivity that meets or exceeds the desires expressed in the NMCI contract award [MTX02] [PEOIT00]. Looking further into the statistics of the Matrix website we can see this level of performance holds true for the availability performance of global internet [MTX02]. Assuming then that NMCI will provide availability on par with that of existing large ISP's, it is logical then that the connectivity of NMCI is adequate for this mission area.

2. Security

The security architecture presented by NMCI is arguably superior to the average defensive measures provided by commercial ISP's. NMCI is held to significantly higher standards prescribed by National Security Agency and the Department of Defense to maintain secure communications [RA01]. The IDS and firewall configurations are likely to be best of breed within the industry when fully deployed. The multiple enclave system established within NMCI is more extensive and comprehensive than common business applications of security. In addition, NMCI makes extensive use of link encryption within the network for long-haul communications which is uncommon in the private sector. Additionally, the level of encryption used within NMCI is of greater sophistication and complexity than presently available in the commercial sector and is the equivalent of that used for secure voice communications by DoD. This is not to say that the NMCI system would be impervious to attack or penetration, but the level of encryption and coordinated defense is much

less interdependent than that of other enterprises operating in the internet [CERT02]. The use of host configuration monitoring and active network management, coupled with layered defensive systems, make NMCI as defended if not more so than the commercial networks or ISPs. For these reasons it is assumed that the security of NMCI is adequate for the conduct of the Administration mission area.

3. Quality of Service

The Administration mission area makes no requirement for the network to handle or segregate traffic internally based on any priority status. This mission area can be accomplished using best effort service as there is no requirement for arbitration for use by the network members. A reduction of throughput will affect all members equally, sharing the adverse consequences that result in an equitable fashion. The quality of service demand for the Administration mission area is no greater than that needed to provide basic internet connectivity in the public domain. Millions of businesses employ the internet for the conduct of normal business today without the need for a specific or on- demand quality of service requirements. The best effort quality of service offered by NMCI is then adequate for the Administration mission area.

F. FORCE PROJECTION

For the purposes of this examination, force projection is defined in the terms of the functions necessary to train, ready, provision, position, and then deploy combat forces to a theater of operation for either combat operations or presence in support of the national security objectives. Force projection differs from administration

in that the actions are directly related to the deployment of forces and all logistical support required placing them in a combat ready status. Movement of personnel for force projection is typically, but not exclusive to unit level movement.

The transactional requirements placing a unit in deployable or combat ready status are influenced by the FAD or Force Activity Designator. The FAD is the mechanism that establishes the logistical support priority to the unit as it approaches its deployment date and then departs for the theater of operation [15]. Once the unit has received a FAD equivalent to deploying status, their transactions, either personnel, material, or movement, would fall under the Force Projection mission area. By the same standard, units that are deployed for presence or actual combat would fall under this definition of Force Projection.

Units that are actually deployed are considered under this definition as many of their logistical requirements are handled by or from organizations that reside within the continental United States. By virtue of this their traffic would require a position in the hierarchy of NMCI to receive the appropriate handling. Force Projection functions reside primarily in the preparation, positioning or provisioning of forces before or after an engagement.

1. Availability

The force projection mission area requires a very high degree of availability, but not necessarily a guaranteed 100 percent availability of the network to be successful. Under the NMCI contract award, SLA 10 provisions NMCI Intranet availability at only 99.8 %. This is significant

in that it is less than the 99.99% that could be provided by simple dual threading. However SLA 24 does provision WAN availability at 99.99%, giving a very high degree of connectivity within a particular Naval Region. Preparations in the deployment cycle possess predictability about them based on the long term planning and coordination required to execute such deployments, much of which is done within the units own Naval Region. This lead time and regional planning function mitigates the lower availability provisioned for NMCI Intranet performance.

It should also be considered that there already exist multiple paths external to NMCI that in essence provide a fail-over protection that could preclude a denial of service to the deploying unit. The Defense Messaging System and commercial or government dedicated voice circuits can provide an adequate means of communicating with any unit or organization supporting the Force Projection mission area. These systems, while supported by NMCI are not wholly contained within the network and should be available should a failure, compromise or denial of service of NMCI occur. Employment of these systems is a risk management strategy for the Navy and provides a means of contacting units in the event of an emergent requirement to deploy.

Should a crisis situation occur that requires the deployment of assets ahead of schedule or on an unscheduled basis the orders would likely to fall to units that are in the early stages of the deployment cycle. Response time of the affected units would be in terms of days or weeks. Other Navy owned communication systems could provide adequate connectivity if NMCI were unavailable.

The net affect of alternate paths provided by other Navy systems, the greater provisioned reliability of the WAN, and the long lead planning for deployments is to mitigate the lower availability of the NMCI Intranet and make the availability of NMCI as an enterprise system adequate for the force projection mission area.

2. Security

The security requirements under the force projection mission are more significant than that of Administration because of their significance to national security. Failure of security at this level represents a loss of strategic advantage to the United States by allowing an adversary to more accurately estimate the response time or capability of American forces to a threat. Opponents of the United States could gain valuable insight into the level of readiness of Navy and Marine Corps forces if they were able to compromise data within NMCI relevant to the Force Projection mission area. This constitutes a greater requirement for NMCI security than under the Administration mission area.

Much of the readiness and disposition of forces data required for Force Projection would fall at or below the SECRET/NOFORN classification level. The SIPRnet is certified to carry information classified up to and including SECRET/NOFORN. The SIPRnet will interface with and is supported by NMCI and will become an integral part of NMCI once security certification is achieved for NMCI. The mechanisms for transporting classified data within the SIPRnet are being applied to NMCI. Specifically the use of link encryption for the transport of data over long haul lines between the NMCI networks operating centers (NOCs).

Certification for NMCI to perform this function is still pending.

Certification for NMCI can be viewed as problematic. There may be significant problems to overcome, however those would largely be technical in nature and will ultimately produce a binary result, certification or non certification. A failure to certify NMCI for the transport of classified information would cause NMCI to be unsuitable for much of the force projection mission area. If however, certification is achieved, NMCI security would be adequate for the performance of this mission. The suitability of NMCI to perform this mission area is then conditional upon NMCI receiving certification for the transport of classified data.

3. Quality of Service

The Force Projection mission area will impose some requirements for differential service within NMCI. There will be a need to give the higher priority mission function (Force Projection) a higher degree of service within NMCI. This requirement is driven primarily by the need to communicate during periods where there is limited throughput due to a compromise or damage. During these periods it should be the priority to forward traffic that is most relevant to the primary mission of the Navy.

Currently, the primary mission of the Navy is forward presence and consequently then the highest priority traffic within NMCI should be that generated by or for the support of deployed or deploying units. The question then becomes how to discern which units have priority. The simplest discriminator between units that is presently available is the assigned units FAD. The FAD is a logistical

discriminator that provides a prioritization to requests from deploying or deployed units. Segregating traffic on the basis of FAD assignment will give the deployed units, or those preparing to deploy, a higher priority cueing within the network.

At present, the NMCI system doesn't employ any specific method for the prioritization of network traffic based on a hierarchy of needs or mission. NMCI is constructed to provide essentially best effort service. The question then is how to provision such a prioritization scheme within NMCI that would fit into the existing structure. There are essentially two different models for provisioning quality of service within the confines of the existing NMCI network protocols. They are Integrated Services and Differentiated Services.

The Integrated Service model is based upon the reservation of service along a negotiated path for the transport of the data [SRSD99]. Before any data sent the resources necessary and the path to the destination are determined and reserved. The weakness of this method is that a path is negotiated before the transmission and so creates a single point of failure. A failure along the route would cause the transmission to be lost or cause the entire process to be reinitiated. When viewed in the context of network survivability, a protocol that creates a single point of failure is less than desirable. This process also fails to make full advantage of the multi-path capability of a networked system. Given the need for a quality of service implementation is greatest when the network is under duress, a protocol that reserves services for a single path is unlikely to provide the desired service for a survivable network.

The Differential Services model has as its premise the differential classification of packets for their prioritization for movement through the network [SRSD99]. The prioritization scheme as applied in the commercial environment is complex and requires ISPs to negotiate bandwidth allocations based on the economics of the services provided. For there to be a consistent application of services the cooperating enterprises must have equitable agreements for the application of bandwidth. If however, the service is to be provided within a single domain, it would be a simpler proposition.

Applying a quality of service scheme within NMCI may be possible by employing a differential services method. The quality of service granularity would be rather coarse, but could provide a means of arbitrating traffic flow during times of high bandwidth demand or low availability. Traffic could be prioritized first by FAD (deployed vs. non-deployed) and second by classifications: routine, priority, or mission essential, in ascending order of precedence. The routine and priority classifications would be the equivalent of what the naval message system currently uses. The mission critical classification could be on the immediate or flash precedence. The ability to assign a traffic classification could be awarded to users within the local area network by the network administrator in a manner similar to that used for the release of naval message traffic. This places the control and responsibility at the local command level where it can be best managed and controlled.

Appendix B contains a more detailed discussion of how quality of service could be implemented within NMCI, but there is the potential for a scalable solution that could

segregate traffic based on a very modest hierarchy. The net result would be to produce a quality of service adequate for the Force Projection mission area.

G. BATTLE MANAGEMENT

Battle management is a relevant mission area for NMCI in spite of the statements limiting NMCI to the confines of the continental United States. In the light of the terrorist attacks of September 11, 2001 it is quite conceivable that NMCI may possess a critical role in the mission of homeland security.

NMCI as designed is part of the integrated network defined as the Global Information Grid. The implication is that to be functional the systems must be compatible. To be interoperable the system must have equitable capability internalized in their design. Whatever can and must be done in one system can and must be done in the others for the network to be successful. Failing to ensure either of these means we no longer have an information grid but an "information island chain". While supportive of one another they are essentially independent because of the lack of similar capability.

The growth of networks has emphasized both compatibility and interoperability to the point of creating what are essentially large homogenous cooperative networks that constitute the World Wide Web. To achieve an information grid, the existing IT-21, NMCI, and other service networks will need to merge at some level. It is logical that all these systems be both compatible and interoperable. Therefore it is logical that whatever Battle Management functions are implemented in the afloat force be supported by the terrestrial WANs (NMCI).

Implementation may not be simultaneous, but is likely inevitable and should be considered.

For the purposes of this examination, battle management is defined as the active command, control and communication with field level units for the purpose of conducting combat operations against an enemy force or the coordination of emergency services (fire, police, etc.) within CONUS in support of the homeland security mission. The relevant activity for the network then is how it facilitates the communication to the field units and or the control of weapon systems used in the actual combat. To perform this mission effectively the network must communicate on a near real time or real time basis. Networks must also provide a high degree of integrity and confidentiality for the data to be of valuable to the network members. Lastly, for the network to reflect the mission priorities in the battle management scenario, it must possess a granular quality of service that is flexible to the needs of the commander managing the conflict. Priority assignment within the network must be responsive to the demand of the mission commander so that the priority mission is receiving the priority service.

1. Availability

As a means of communication to both internal and external organizations an IP based network has a well established history. The flow of information between individuals has become comparatively reliable as evidenced by statistics available in the public domain on network performance. Based on the 26 internet service providers surveyed by the Matrix.Net site discussed earlier under Administration, 20 met or exceeded the standard performance

metrics (SPM) for NMCI in terms of availability, packet loss and average latency [MTX02]. The question then is whether the SPMs within the NMCI contract award adequate for the mission. The relevant SLAs and their standard performance metrics (SPM) for network availability within NMCI are summarized in the table below.

Service Level Agreement (SLA)	Availability Requirement in NMCI Contract Award
SLA 6: Web Access Services	99.5%
SAL 10: NMCI Intranet Performance	99.8%
SLA 11: NIPRNET Access	99.5%
SLA 12: Internet Access	98.0%
SLA 13: Mainframe Services Access	99.5%
SLA 14: Desktop Access to Government Applications	99.5%
SLA 18: Unclassified Remote Access	99.5%
SLA 19: Classified (Secure) Remote Access	99.5%
SLA 24: WAN Network Connectivity	99.99%
SLA 25: BAN/LAN Communications Services	99.99%
SLA 35: Information Assurance Operational Services - SIPRNET	98.0%

Table 7 NMCI Contractual SLA Availability Levels

(Summary of SLAs taken from the NMCI Contract Award of 6 October 200, NMCI Contract N00024-00-D-6000)

Not all of the service requirements are provisioned to meet the availability threshold of a dual threaded system as discussed in appendix A. Only SLAs 24 and 25 are provisioned to the level of "four nines" of availability (99.99%). These SLAs refer to the availability or connectivity of the wide area (WAN), base area (BAN) or local area (LAN) within NMCI. The consequence of this is

that the highest levels of connectivity are assured only at the WAN or below. These small differences in the percentage availability are significant because they represent the difference between a system that possesses redundant capability and one that contains a single point of failure.

Provisioning availability in this manner creates islands of higher availability within NMCI. The true requirement for availability in this mission area is dependent upon how this mission function is accomplished organizationally. If the battle management function is ceded to a particular region, the WAN availability may be adequate. If not, then the lower availability of the NMCI Intranet (SLA 10) is likely inadequate based on the less than dual threaded level of availability that has been provisioned.

Without provisioning a higher degree of availability across the entirety of NMCI, particularly the very high speed backbone network system (vBNS) that connects the NMCI WANs nationally, it is unlikely NMCI is suitable for the Battle Management mission area.

2. Security

A network employed for battle management places a premium on the demand for integrity and authenticity. Even when data is received in a timely manner there must exist a very high level of assurance that (1) what was received was what was sent and (2) the sender is exactly who the receiver thinks he is. If either of these conditions cannot be met the value of the data is seriously in doubt. High degrees of authenticity and integrity can be achieved in a number of ways in different networked systems. Recall

the example Alpha-Xray Battle Group Communications example used earlier in the chapter. A similar lesson on integrity and authenticity can be drawn from examining carrier air wing strike operations.

It is common practice for air wings to conduct strike operations using clear voice (non-encrypted) channels. This practice has developed because of the difficulty in achieving full connectivity among all aircraft using encrypted voice transmissions. The link encryption systems can be problematic when implementing across a large number of aircraft. The air wing has been successful at this because of the trust that is built within every layer of the network (air wing, squadron, flight, and section) during pre-deployment training. The individual nodes become very familiar with the other nodes with which they routinely interact, to the point of being able to recognize the other's transmissions.

This trust relationship is developed over the multiple rehearsals in the primary mission area of the network (air wing). The network is decentralized and mission responsibilities are resolved down to every node within the network. Data integrity is checked by the receiving node through the validation of the sending node (does the pilot recognize the voice/call sign) and by logical comparison of the message with the environment as the receiving node understands it currently and has experienced it in the past (does the message make sense when referenced to the pilots situational awareness).

This method is not unlike the pattern recognition and anomalous behavior used by intrusion detection systems. The difference is that IDS are used in IP networks to prevent denial of service attacks while this voice

recognition and nodal behavior patterns in voice circuits are used for integrity and authenticity. The key is that the voice circuits have many intrinsic clues that are unavailable to the nodes of a computer network. The sound of the voice, the background noise, the clarity of transmissions all play a role in helping the aircrew authenticate the source as trusted or untrusted. The aircrew (nodes) develops a historical log for evaluating the integrity of the data provided by the sender and sometimes a factor of error correction applied to achieve a common picture. IP networks do not have the luxury of such interpretation and therefore require a more binary measure.

In the IP world, the digital signature is the analogous system to voice recognition used in the air wing clear voice communications. PKI is a suitable solution; however the scale of the problem is immense. The number of certificates to be issued and managed for even a small battle problem would be significant considering the number of ships, aircraft, sensors, and weapon systems that would feed a common battle management network. Some form of object level security is likely the answer to this problem. Object level security could give the individual nodes the authenticity and integrity capability that equates the voice recognition capabilities of radio networks. The difficulty comes in the full implementation of the PKI system within NMCI and across the Navy as an enterprise. This is a large and formidable problem that has yet to be completely resolved. The Force Net project under development now may answer many of these issues and its emphasis is primarily on the cohesion of shipboard combat systems first, with the integration of terrestrial networks coming later. Given the emphasis on shipboard systems in

the Force Net project and the lack of full implementation of PKI within NMCI, the security of NMCI is likely inadequate for use in the Battle Management mission area. Until PKI or other implementations can be deployed and their authenticity, integrity, or security validated it is doubtful that NMCI could support the Battle Management mission area.

3. Quality of Service

The quality of service demands for the battle management mission area would be as complex as the requirements for security. In addition, any quality of service scheme for battle management would also have to be extremely flexible. A quality of service management scheme would need to be responsive to commander's intent on a real time basis. This is driven by the need of the commander to emphasize a particular geographical area (over land, over water, inner zone, outer zone) or warfare specialty (anti-submarine, anti-air) based on the developing engagement. This would guarantee the traffic deemed most relevant to the mission commander gets the priority handling to and from the end nodes. The routers handling this information must respond to the prioritization, or reprioritization immediately for the network to be an effective system for managing an engagement that captures all the required data. If the scope of the data handled by the network was limited, a method may be found to achieve some prioritization. It is possible that, should the data flow be small enough, and the throughput large enough, quality of service may become irrelevant. In this situation best effort service may be adequate for most transmissions.

Ultimately, the requirement for quality of service for battle management is driven by the system or systems with the highest quality of service need. If any single application requires guaranteed service, then the only way to presently implement such a scheme (integrated service) creates single points of failure for that application. If the demanding application is pivotal to the battle group defense then an opponent need only destroy a single node to deny the entire battle group the advantage gained by having the network in the first place. The answer is complex in that to resolve it there must be a resolution between competing demands for bandwidth "on demand" and the delivery of guaranteed service to specific systems. NMCI will service a deployed Fleet that will be reached via satellite or radio WANs that presently experience limited bandwidth availability. Until these issues are resolved it is doubtful NMCI will be suitable for the Battle Management mission area. NMCI as it is being deployed now is likely unsuitable for this mission.

H. MISSION AREA SUMMARY

Reviewing the mission summary for NMCI we can see there are shortfalls in what the system needs to be to complete all these mission areas effectively and what is. The question resolves to what mission NMCI could do best now. The obvious choice from the perspective of the existing network capability is Force Projection. The table below summarizes the evaluation.

Requirements			
Mission Area	Availability (Ao)	Security	Quality of Service (QoS)
Administration	Yes	Yes	Yes "Best Effort"
Force Projection	Yes	No <i>Pending¹</i>	No <i>Possible Solution²</i>
Battle Management	No	No	No

Table 8 Mission Requirements Summary Matrix

1. Pending certification will resolve this to a "yes" answer

2. A scalable solution exists to produce some level of QoS within NMCI, but is not implemented

The Administration mission area is easily suitable as this mission area entails the deployment of an enterprise wide network. NMCI is this very thing for the Navy. The mission requirements for this fall completely within the intent of NMCI from its very beginning and are fully supported in the implementation.

Battle management is a demanding mission for which NMCI, or as yet any other IP network, may not yet be ready. The FORCENet program, under development by the Naval Warfare Development Center (NWDC), Space and Air Warfare Systems Command (SPAWAR), Naval Air Systems Command, and the Naval Sea Systems Command is an effort to congeal all of the existing tactical networks into a single coherent network within the shipboard units. Once that is achieved the requirements to implement the Battle Management solutions within NMCI must be balanced with any associated costs and potential gains from it. There will also need to be a discussion and decision as to whether this mission

function for a terrestrial WAN is appropriate tactically or organizationally within the Navy.

The majority of the functions to support the force projection mission are in place within NMCI with the exception of quality of service. The availability needs in the force projection mission area are mitigated by the designed lead times and semi-routine nature of the mission in its execution. Security certification is likely to be approved and the system deemed capable of supporting the security requirements. A quality of service application within NMCI is possible using existing technologies and while coarse, would provide a means of segregating traffic.

The significance of resolving the quality of service problem should not be underestimated. Implementing quality of service within the network is what provides it with the self sealing qualities it needs to become a survivable system. Security and availability are important, but there are viable existing solutions that could be applied to those issues within NMCI. Quality of service can be viewed as the long pole in the NMCI tent. Even if security is adequate, availability is high, and the mission is well defined, there must be a quality of service upon which the network can rely to allocate the available bandwidth. Security, availability, and mission definition serve to create a connection between the nodes, and predetermine a hierarchy of functions within the network, but quality of service makes the decisions of how to use the bandwidth that is made available.

Examined within a purely naval frame of reference, security, availability, and mission are the damage control parties for the network, keeping it afloat when damaged by an attack. Quality of service is the network captain,

determining how to fight the ship and complete its mission with what is left unscathed.

I. THE NEW MISSION DEFINITION FOR NMCI

"... amateurs talk tactics, professionals talk logistics..."

Having examined the potential major mission areas for NMCI in the previous pages, we can now redefine the mission for the Navy and Marine Corps Intranet. The previous discussion showed that the Force Projection mission area is the most applicable to NMCI as it is being deployed. Redefining the mission of NMCI as Force Projection will accomplish three very important things for the Navy. First, it aligns the central mission of NMCI with a core mission of the Navy and Marine Corps team.

Second, it should provide near and long term economic advantages to the Navy by the disintermediation of the supply chain, placing the combat units (consumers) closer to their suppliers (the supply system) and shifting more power to the consumer organizations just as the internet created "reverse markets" in the business world [JH97]. And lastly, placing the logistical and readiness functions within NMCI will infuse it (NMCI) within the modern Navy culture and begin to harness the great potential network centric warfare may offer. The question in the context of survivability then is to define the mission essential functions NMCI will need to perform in this mission area to provide those benefits.

1. Mission Essential Functions

The functions of NMCI under the Force Projection mission can, on an aggregate scale, be segmented into two primary functions with two sub-functions under each. Table 2-2 shows the concept of how these mission essential functions could be logically divided. In the fullest application of the Survivable Networks System Analysis method the mission essential functions should be identified specifically to the network, but under this examination and given the vast scope of this problem such a process is impractical within this thesis. However they can, on an aggregate level, be logically subdivided in this manner and the matrix used as a guide for the process of identifying and prioritizing those mission essential functions.

Mission Essential Function Categories	Support of Deployed Forces		Support of Non-Deployed Forces	
Mission Essential Function Sub-Categories	Logistics	Readiness	Logistics	Readiness
Mission Essential Function Flows	1. Units actively engaged. 2. Units preparing for engagement or redeploying 3. Units providing presence	1. Units actively engaged. 2. Units preparing for engagement or redeploying 3. Units providing presence	1. Units within 90 days of deployment. 2. Units ending/beginning deployment cycle	1. Units within 90 days of deployment. 2. Units ending/beginning deployment cycle

Table 9 Force Projection Mission Essential Functions

The primary functions can be categorized along the lines of deployed and non-deployed operations. Deployed operations are those functions necessary to support deployed units or those units with 30 days of departure. This roughly corresponds to the FAD designation mechanism that exists within the logistical system today and provides a division of emphasis that aligns with the operational concepts of the Navy and Marine Corps. Non-deployed operations are then the supporting functions required for units that are in garrison or home port and are outside of 30 days of departure in their deployment cycle.

The mission essential function sub-categories are logistics and readiness. The logical division here is based upon the high level criteria used by commanders in their decision process in their selection of units for any particular operation. Operational commanders need to understand the logistical and operational readiness status of any combat unit in their area, what their strengths or deficiencies are and when they are likely to be resolved. Supporting commanders must have visibility of the continual progression of a unit's logistical and operational readiness status throughout the pre-deployment cycle. This allows the supporting command to emphasize the training and logistical requirements that are relevant to the individual unit and align them with the needs of the operational command that will gain the combat unit once it is deployed. These functions occur now, but not in any networked environment and not on any real time or near real time basis as would be capable under NMCI. These two mission essential function sub-categories embody the Force

Projection mission, which is readying, deploying, and then supporting combat units around the globe.

Beneath each mission essential function sub-categories there are identified mission essential function data flows. These are meant to classify and prioritize the primary data flows within each mission essential function sub-categories. The emphasis again is on the operational requirements of the Navy and Marine Corps team. An operational commander's concern lies in the logistical and operational readiness of his engaged units first. Next is the readiness of his disengaged or redeploying units and what their needs are to return to full combat strength. And third he must have knowledge of the logistical and operational readiness of his available supporting units. The supporting commander's requirements are the basically the same. He must understand the status of his units in the pre-deployment cycle in order to provide the operational commander the combat units when and where he needs them and trained for the appropriate mission. The net result of structuring the NMCI mission functions in this way produces gains that can be realized in both deployed and non-deployed operations.

a. Logistics

Operational commanders will gain from the near real time data provided by logistical data flows that would be available from this mission realignment. Battle groups and Marine Expeditionary units currently deploy with their own networks that are linked to CONUS via satellite or other means. The development of common data bases linked via the existing networks could feed theater commanders and Fleet CINCs important logistical data. This data is

presently transmitted via the naval message system and a multiple of individuals and systems are required to collate and respond to those requests. An integrated logistical data base riding on NMCI would minimize the operational units required inputs while increasing their visibility within the operational chain and the supply system. At the same time, operational units would obtain greater visibility into the supply system.

Operational commanders would be able to view requisition status, priorities, and supply status in a real or near real time basis. They can demand more timely response by creating competition for the delivery of the service among internal Navy providers. Individual suppliers will no longer be tied to customers geographically as they are today. If the required components can be located, the only requirement is to coordinate the delivery. This method is more like the just in time inventory methods employed by businesses today.

Non-deployed commanders will also gain from the increased visibility within the supply system but the primary benefit to them will be monetary. The gain within non-deployed operations will be in the efficiencies created by the disintermediation of the supply system. The greater visibility up and down the supply chain will again create competitive pressure among the members of the supply system. The result will be to eliminate the inefficiencies that exist under the hierarchical supply system of today. The resulting supply system will likely represent more of a mesh and less of a chain. The net result for deployed and non-deployed units is greater supply system responsiveness at a lower cost to the operational units.

b. Readiness

Operational commanders will benefit by greater visibility of the readiness status of their forces. This data is presently collected via the Status of Readiness and Training System (SORTS) reporting. The data is provided by the individual units but is not real time or even near real time. It is taken on face value that the status is valid unless otherwise reported, but in times of crisis operational commanders may not have the time to revalidate the data. A near-real-time system would provide the answers quickly and allow those units requiring additional or refresher training to have those needs identified in a timely manner. Supporting commands can view the specific requirements of the gaining commands and train and prepare the deploying combat units accordingly. This is particularly useful when operational requirements change or supporting unit are to be provide by the Reserve Force organization. The supporting command can anticipate the gaining commands needs and plan accordingly. Greater visibility will require more accurate reporting by the combat units, but ultimately a more clearly accurate assessment of unit readiness can be made.

J. CONCLUSION

It is well understood that there are significant challenges in the implementation of this mission area. There are many applications and data bases that must be merged into common forms for all of this the functionality to emerge and provide the data flow necessary to make the logistical and readiness benefits achievable. This transition process is, however, an implicit part of the implementation of any enterprise wide network such as NMCI.

Many businesses have engaged in this transition successfully, though possibly not on the scale faced by NMCI.

The significance of the mission redefinition for NMCI is that it will focus the efforts of that transition process in a way so that when it is completed, the NMCI structure and function directly supports the core mission of the Navy and Marine Corps. The present transition to NMCI is driven by the need to connect all the nodes, without real a great deal of consideration as to what the central mission of the network should be in support of the Navy and Marine Corps team.

The projection of combat force around the globe in support of U.S. national interests is a core mission of the United States Navy and Marine Corps and by logical extension so should it be for NMCI. Aligning the mission of NMCI with the operational mission of the Navy will ensure NMCI becomes a viable part of the operational capability and will provide the greatest benefit to the individuals at the tip of the spear.

III. LEGACY SYSTEMS AND NMCI

A. WHAT IS LEGACY AND WHY IS IT IMPORTANT?

The term legacy is used extensively with the IT community, but it does not necessarily mean the same thing to everyone. One man's legacy 386 computer is another man's upgrade. In a very general way, the term legacy describes the existing components (hardware and software) that constitute a network (local or wide area) during a transition to a newer network standard. The terms "legacy" and "upgrade" are moving targets, but they define one another during a fixed period of time. Therefore, the definition of what legacy is or is not must be based on a common reference.

NMCI, like many enterprise networks endeavors must have a fixed reference since it is an inherently moving target. The definition of legacy in terms of NMCI will change as NMCI becomes an ongoing concern for the Navy and Marine Corps. Iterations of the NMCI contract could bring a new reference for what is legacy and what is not. In essence then, legacy can then be defined in terms of compliance, or the lack there of, of existing application and network standards established by the most recent NMCI contract award.

Under this examination then, the reference would be the 6 October 2000 NMCI contract award by the DoN PEO-IT. The advantage for the Navy and Marine Corps is that those systems under the NMCI umbrella will be brought into compliance with the NMCI standards as part of the NMCI contract. The disadvantage is that for an undetermined period of time, legacy applications and or legacy networks

will operate in the NMCI environment in non-compliance with the established standards

Legacy networks and applications within the Navy and Marine Corps are one of the specific targets of the transition to the NMCI environment. NMCI was contracted as a service to address the specific problem of multiple stove-piped networks that exist within the Navy and Marine organizations. Procurement of the necessary hardware and software for the operation of local area networks was the responsibility of the organizational level authority in the pre-NMCI environment. Under NMCI, this is no longer the case. The great success of the NMCI contract to date has been the consolidation of the many small organizational level networks under the NMCI umbrella. Doing so has enhanced the overall security of the organizational and higher echelon level networks by producing a largely consistent suite of hardware and software that are compliant with the established security and interoperability standards.

In addition, the Navy has achieved cost savings by eliminating the cost of operating and maintaining these smaller nets. There are, however, significant Navy and Marine legacy networks and applications that can not or will not be transitioned under the initial NMCI contract award. The largest and most relevant of these networks are the Marine Corps Enterprise Network (MCEN) and the IT-21 standard shipboard networks [RAY01].

The significance of the inclusion of Marine Corps Enterprise Network and the IT-21 legacy networks is that their presence is necessary for the larger success of NMCI, while their noncompliance to the standards within NMCI present a vulnerability that could be exploited. These

existing, older systems create a tension between what is best for NMCI network security and what is best for the organizational and mission functions of the Navy and Marine Corps and NMCI. The need to include these legacy networks and applications is driven by the operational and fiscal realities faced by the Navy and Marine Corps. For NMCI to be effective at its mission it must include these older networks and the information and connectivity gained from them.

If NMCI is to be a part of the greater global information grid, it will certainly be required to interface with dissimilar network structures that may or may not meet the same security and application standards. These networks could be intra-service (MCEN and IT-21) or inter-service in origin. Creating NMCI as a bounded (closed) system would improve its security, but would ignore the mission functions of the other external networks and their relevance to the overall NMCI mission. Doing so also ignores the fiscal realities of the investment made in those external networks and the logical goal to eventually merge NMCI with the other Navy and Marine networks into a single coherent system **[RM02]**.

So, to be effective both now and in the future, NMCI must be able to accommodate the existence of legacy networks and applications in some manner and for some period of time for it to be mission effective. For this reason, the effective transition of legacy applications and networks is relevant to NMCI now and in the years to come. The focus should be on how to transition legacy without making diminishing the positive effects of new standards for security and availability. Legacy transition for NMCI during the initial cutover is of particular importance

because it will lay the ground work for the future success or failure of the network.

B. LEGACY TRANSITION

Legacy systems represent a distinct challenge for any enterprise network implementation either in government or the private sector. While there is no specific discussion of legacy transition under the survivable network system analysis method, legacy networks and applications are relevant to NMCI network survivability in two ways. First, the requirement to operate with legacy inter-service and intra-service networks places NMCI squarely in the realm of unbounded networks as discussed by the authors of the NSA model.

NMCI will not be able to impose any administrative control or likely possess great visibility within these other external networks (domains). This is effectively true now with NMCI interoperation with the MCEN and IT-21 which are defined as legacy networks under the initial NMCI contract award. The IT-21 in particular and MCEN (to a lesser degree) networks are comprised of nodes that enter and exit the NMCI environment at will. These deployable nodes (USN ships and USMC units) operate independently with other external networks and then return at their own will and reconnect to NMCI. There is an interface between the IT-21 node (ship board LAN) and NMCI, but the IT-21 node is still directly connected to NMCI. The same can be said of the deployable NMCI seats.

These hardware items are contractually procured for the purpose of moving between the NMCI and IT-21 network for the purpose of supporting the deployable staffs and aviation units. NMCI must interface with these networks to

be mission effective, and so by necessity rely on a trust relationship with them to function in a secure manner. This necessity effectively makes NMCI an unbounded network and therefore the concept of survivability is totally relevant to NMCI.

Second, the size and scope of the legacy transition for NMCI presents a tremendous challenge to the security of NMCI. At present there are approximately 37,000 legacy applications and 400 legacy terrestrial networks that exist within the Navy and Marine Corps organization [GCN02]. Many of these legacy systems have been identified as mission essential by the organizations that employ them and are considered integral to their function and continued operation. The inclusion of these and other legacy networks undermines the "hard target" approach that has been taken in the design and implementation of NMCI security.

1. Legacy Networks

Recall that the NMCI security architecture is enclave-based, employing the concept of defense-in-depth at each using the same means of perimeter defense at each layer. The inclusion of these legacy networks produces potential soft spots in the defense in depth concept by taking software and hardware that is known or assumed to possess security vulnerabilities (by its definition as a legacy system) and attaching it to a more hardened network through valid interfaces. The net effect is to present an intruder the opportunity to compromise the NMCI environment via legitimate channels after penetrating an assumed weaker defensive perimeter. At the end of the day the net effectiveness of the defensive mechanisms for NMCI are no

better than the weakest link in the perimeter. In the case of NMCI that weakest link is the attached legacy network or application. Attaching NMCI to these other legacy networks weakens the enclave strategy and hence the security of the entire NMCI network.

The problem for NMCI is that it must operate with these other networks to be mission effective in spite of their potential weaknesses. Consequently then, additional measures should be taken to mitigate the weakening of the overall network security by legacy participation with NMCI. The network survivability analysis method can be applied to NMCI and the transition of legacy to assist with this problem.

2. Legacy Applications

At present there are 37,000 identified legacy applications that must be dealt with by the NMCI ISF. While it is unlikely that all of these applications will be transitioned into NMCI, certainly a portion of them will be required. The requirement will be driven by the mission essential functions that the legacy application provides to the organization. Some systems, regardless of their age and format, will not or cannot be transitioned in a timely manner. The consequence is then that some of these will persist within or attached to NMCI in some manner. For the Navy and Marine organizations to benefit there must be the greatest achievable range of connectivity. This requirement for connectivity to these "mission essential" applications creates vulnerabilities in the same manner as the inclusion of the legacy networks, which may contain some of these legacy applications.

The inclusion of these legacy applications undermines the hard-target approach to the enclave security architecture. Building a new network that includes older software that does not meet the newer security standards reduces the overall effectiveness of the integrated defense. The implications and net effect for these legacy applications on network security is the same as the inclusion of legacy networks. The newer standard will be only as good as the legacy applications.

An alternative to the inclusion of legacy applications within NMCI would be to host them separately in a separate networked environment. A "quarantine" strategy is being examined for the hosting of legacy applications that are considered too insecure or too difficult to make compatible for inclusion with NMCI.

The quarantine concept is based upon the construction of a shadow network, external to NMCI, which provides users with access to those mission essential applications that cannot be hosted internally to NMCI. The quarantine method, while providing access, runs counter to two of the basic tenets of the NMCI contract, the elimination of non-standard networks and the achievement of total interoperability. The NMCI services contract was intended to and has performed well at consolidating the disparate networks operated by organizational level units.

Producing a shadow network to host these legacy applications reverses that successful trend. In addition, the shadow network produces a security risk at some level. The thin client machines required to access the legacy applications would undoubtedly be co-located with the other NMCI hosts. Reliance on an "air gap" between the two

networks may appear to be adequate insurance against compromise, but in the end it is still exploitable.

Social engineering attacks are commonly used to avoid the internal security measures of contemporary networks and there is no reason to believe that this tactic could not be successfully employed against the shadow networks and then ultimately NMCI [TR02]. In the effort to accommodate users the undercurrent of legacy applications could re-grow the networks eliminated by the services contract while creating a security weakness that has very little visibility.

This is completely contrary to the two primary goals for the NMCI services contract which are the reduction of stove-piped networks and cost savings. Even if the security implications are completely ignored it is difficult to accept the obvious reduction in dollar savings to the Navy and Marine Corps as a result of this option of legacy inclusion. Given the Congressional oversight of NMCI to date, creating a significant funding requirement for a shadow network over the long or short term could endanger the existence of the network more effectively than any potential intrusion.

There are still other alternative strategies for legacy applications that can be considered. All of the legacy applications could be excluded until they are "de-loused" of problems or validated to meet security requirements. While this method places greater emphasis on security, it assumes that all legacy applications are of equal importance. Unless the relevance of the individual application to the core NMCI mission functions is considered, there is the likelihood that the selection and sequence of applications for transition will be haphazard, inefficient or ineffective when compared to what the

mission essential functions are within NMCI. The result is that the emphasis of the transition could then be less than optimal, producing greater costs and more time delays than would otherwise be necessary.

Alternatively all legacy applications could be included in NMCI in a "Mariel Boat Lift" manner, bringing all of them into NMCI for convenience and then attempting to find the bad apples as you go, de-lousing them inside NMCI. This is an equally inefficient and far more risky strategy. Assuming that all the applications could function within the NMCI environment, which arguable they all may not, the likelihood of there being significant exploitable vulnerabilities in these applications is quite high. Accepting such a tremendous risk for a network that has been deemed mission essential by the CNO seems very unsound. Again, ignoring the security implications, this method of dealing with legacy applications does nothing to discourage their continued use and undermines the cost saving strategy of the NMCI contract. There is no leverage on the user to bring his application into compliance with the NMCI security and application standards if they are allowed to continue to function as they did before NMCI. The Mariel boat lift approach is likely too great an accommodation and takes tremendous risks with the security of NMCI, running counter to the business and security models that are at the heart of the enterprise implementation.

Legacy applications and networks are essential to the NMCI transition because they contain valuable information that is relevant to the entire Navy and Marine Corps enterprise as an ongoing concern. The information they contain and the connectivity to the force that they provide

are extremely valuable to the first iteration of the NMCI contract. At some level they require inclusion until a time that these older networks and applications can be brought in line with the desired standards for NMCI. So then, if legacy must be included, the importance should be placed on choosing those legacy networks and applications that directly support the core mission functions of NMCI and the Navy and Marine Corps, and then applying survivability methods to mitigate the inclusion of these less secure networks and applications within NMCI. The survivable network analysis method provides some of the answers to how to organize the transition process and then implement survivability into NMCI. The survivable network analysis method will at least mitigate an inclusionary approach toward legacy systems.

C. NMCI TRANSITION ORGANIZATIONAL METHODS

Before examining the concept of a transition under the SNA method, other more traditional methods for the organization of a transition should be examined. Looking at these methods will provide some background for comparison to the NSA method to determine the advantages of each.

1. Transition by Claimant

This method is focused on the command (or claimant depending on level chosen) and the functions therein that are considered essential for the continued operation of that unit, staff, or other type organization. The intent is to keep the organizational level functions intact and operating while retaining the ability to service the demands of the immediate superior, subordinate, and any

supported commands. This method of transition allows commands to move into the NMCI environment as a whole functional organization. Legacy networks are addressed on an individual basis, being assumed piecemeal as the total unit enters NMCI. Legacy applications are addressed in a parallel manner by separate elements of the NMCI transition team. Each application is categorized by functional area and then the transition team determines if there is a suitable substitute or if the individual application can be transitioned via an application interface [RM02]. The advantage of this method is that it allows units to enter NMCI intact, with all of their core functions operating within NMCI. Unit effectiveness is retained because the organizational functions are consistent across the unit level enterprise. There would likely be no two-tiered arrangement where some command level functions are within NMCI and some without. Unit level costs for support are reduced and unit functional capability may be achieved in a single transitional period. The emphasis in this format is on the prioritization and selection of the Navy and Marine units for transition at each level of the command structure.

The disadvantages to the claimant organizational method is that there is no consideration to mission function of NMCI and the role or the contribution of the individual unit (and its inherent legacy applications) to that mission function or functions. The emphasis is merely on unit cohesion and communication with superior and subordinate units regardless of what the intended purpose of that communication is.

While organizational cohesion is important, retaining the cohesion of mission function is equally important and

could potentially be ignored under this scheme. This method is also iterative at some level. If units are chosen based on claimancy, and legacy applications reviewed in the same manner, the potential exists for discovering a more suitable substitute application after some number of units have already been completed.

The result is that any previous claimant organization that was transitioned may have to transition an application(s) a second time after completing the initial cut over into NMCI. This remedial transition process is inefficient and likely increases cost associated with the transition. The result may be that once the transition of the total force is complete there will be a laundry list of required remedial transition efforts to put the entire enterprise on exactly the same application footing, which was a primary objective of the initial cutover process.

2. Transition by Warfare Specialty

Many of the organizational lines of responsibility within the Navy are drawn along the lines of warfare specialty. This method of transition would provide for a cohesive move of any particular warfare area into NMCI as a community of interest. This method would move larger segments of the Navy organization into NMCI at one time but would require more coordination for the actual cutover. This would delay the specific transition date at the unit level; however the delay would be off-set by the resulting cohesiveness achieved within that community by resolving the community's legacy application transition in aggregate vice on an individual basis. The advantage of this method is that larger communities would transition in as a single event, placing more emphasis on mission function within the

communities and potentially better serving more of the mission function of NMCI and the Navy than a claimant approach. The emphasis in this format then is on the prioritization and selection of warfare specialties for transition.

The disadvantage of a warfare specialty approach to transition is that the majority of the Navy's warfare units (surface ships and submarines) are under the IT-21 umbrella, a program separate unto itself specifically identified as a legacy system by the NMCI contract. IT-21 will likely merge with NMCI in the future, but to remain effective the surface units should retain a common network footing until the essentially terrestrial, non-mobile NMCI network is complete. NMCI was not intended for surface units, but for shore-based organizations.

This approach also ignores the area where the majority of the legacy applications reside, in the support organizations of the Navy and Marine Corps. Transitioning the warfare units first would leave the vast majority of the legacy applications and many legacy networks untouched until the final stages of the transition period. While being a mission oriented approach to transition it is uncertain whether the mission areas addressed are relevant to the mission functions of NMCI. In the end the mission functions transitioned may not be the mission functions supported or desired within the NMCI environment.

3. Transition by Navy Region

Many administrative functions are organized along Navy region lines around the globe. Arbitrary functional lines of responsibility are drawn to separate control between fleet commanders, theater commanders, and regional

commanders. These lines of separation provide for the smooth administrative function and allow the individuals responsible the benefit of understanding where their authority begins and ends for the myriad of administrative functions they must perform. The advantages for this method are achieved in the cohesive and comprehensive nature of the transition by region and the multiple functions contained within them.

Many training and logistical functions occur within a particular Navy Region. Most Navy regions contain large numbers of combat and support organizations that perform services for one another and or deploy and operate together underway. Organizing by Navy region would produce a cohesive subnet within the Navy as each region is transitioned. In addition, the functions contained within Navy regions are similar, so the transition of one region would likely be very similar to another. Though there may be relevant and significant differences, the composition of the major Navy regions is essentially the same. Efficiencies produced by the transition could be realized on a regional scale and the lessons learned applied to the transition of successive regions. The emphasis in this format is on the prioritization and selection of the individual Navy regions for transition.

The disadvantages of transition by Navy region are similar to transition by claimant or warfare specialty. First, choosing a region ignores the significance of mission function within the NMCI network. The advantage is gained when the mission function is distributed across the largest possible number of nodes. Segmenting the transition by region delays the ultimate payoff until all the Navy regions are completed. In addition, there may

likely be a remedial transition process once all the regions are complete to assure that all of the regions are on the same application and security standards. Given that the denominator for the transition in the Navy region the remedial effort could be significant. The additional cost and time to complete the effort while not maximizing the mission effectiveness of NMCI could be significant.

Each of these organizational methods for has been applied across a very wide variety of projects within the Navy and Marine Corps. The shortcoming of these plans is that they ignore the potential mission capabilities and limitations of the NMCI network. In each of these methods, the NMCI mission is a secondary consideration in the implementation and the primary focus is on the needs or requirements of the Navy organizational level that is being transitioned. While any of these methods could, in the long term, transition the legacy applications and networks, none of them attempt to optimize the employment of the NMCI structure and function to the war fighting needs of the Navy and Marine Corps.

D. NMCI LEGACY TRANSITION

It should be noted at this point that the majority, if not the entirety of the discussion pertaining to legacy within NMCI references deals with applications and not networks. It can be assumed this is because the vast majority of the smaller local and base area networks can and will be assumed by the NMCI primary contractor at AOR, and the hardware completely replaced by the cutover. The major legacy networks, MCEN and IT-21 are specifically identified in the NMCI contract and as such it is assumed they will likely be the only remaining legacy networks once

the NMCI transition is complete. This assumption holds true only if the operational functions of the Navy are ignored. There will still be significant assets employed by the LINK 16, LINK 11, CEC, and voice nets.

That said it is quite possible that some other smaller terrestrial legacy nets may continue to exist after cutover is complete. If so, the implications of their continued existence is relevant but can not be specifically explored until such time as the transition is complete. The NMCI guidance on transition is explicitly application oriented and the relevance to the networks can only be inferred from those references.

The Navy Marine Corps Legacy Transition Guide, version 2.1, dated 26 October 2001 is the NMCI Program Management Office plan for the transition of legacy applications into the NMCI environment. The specific requirements and procedures needed to transition and support a legacy application into the NMCI are contained within this document. Examining the process of transition should then demonstrate the focus of the transition plan by the NMCI PMO and how the transition for the Navy is organized. The end to end process for the transition of legacy applications is summarized in the figure below **[LTG01]**.

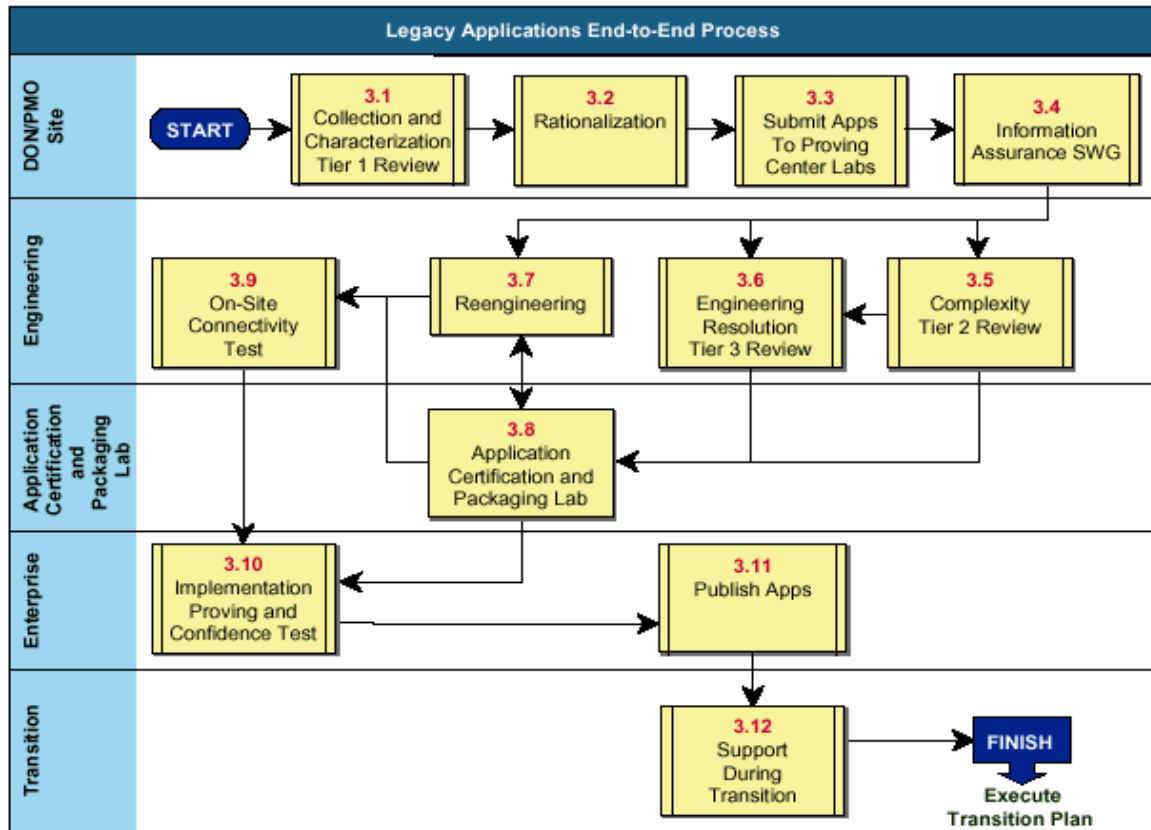


Figure 4 Legacy End to End Process

The transition process is initiated by the individual command as a preparatory to assumption of responsibility (AOR) by the contractor and final cutover into the NMCI environment. AOR is the period when the contractor assumes responsibility of managing the existing Navy IT assets for the Navy organization. Cutover is the point when all the existing Navy IT systems are replaced with contracted hardware and software.

The customer organizations are primarily responsible for identifying and gathering data on the requirements for the transition of legacy applications into NMCI [LTG01]. The generic timeline for this process from beginning to end is estimated at 180 days. The process begins with the collection of data on all legacy applications within the

command using some standardized tools provided by the ISF. It is primarily the customer organizations responsibility to rationalize and prioritize the legacy applications to be transitioned into NMCI. Failure to perform this function properly or failure to identify critical applications in a timely manner will result in those applications not being migrated into NMCI by the cutover date.

The implication of this is that the customer organization must then pay for the specific application to be transitioned. Customers are also encouraged to eliminate redundant, unnecessary, or non-standard applications [LTG01]. This encouragement comes without any metrics for deciding what applications are redundant or unnecessary. The customer must make these initial determinations. When the ISF and PMO personnel arrive on site there are additional data collection and classification of the legacy applications that occurs.

From this point, the responsibilities fall to either the PMO or the ISF, either individually or jointly, to complete the work necessary to transition and certify the applications for operations within NMCI based on the input from the customer organization. Overlaid with this unit level function for transition is the higher level guidance from the DON CIO's office for the elimination of legacy applications.

Appendix E of the Navy Marine Corps Legacy Transition Guide is a Memorandum for Distribution from the DON CIO for the management of Department of the Navy Software Applications dated 23 April 2001. The memorandum directs all of the Navy's Claimant organizations to reduce the number of duplicative, obsolete, and non-secure applications. This process is to be performed by all

Department of the Navy organizations. The memorandum outlines the need to create a structured enterprise IT architecture along functional lines to ensure the horizontal integration of business practices [LTG01].

The DON CIO memorandum references the Navy Marine Corps Legacy Application Transition Guide for the specific processes to be performed for the transition process in support of the Claimants efforts to reduce the number of legacy applications. With these two documents as reference we can step back and look at what the transition plan is in the larger context of NMCI.

E. THE CURRENT PLAN FOR LEGACY TRANSITION

The Navy Marine Corps Legacy Transition Guide and the included DON CIO memorandum of 23 April 2001 combine to show the organizational plan and emphasis for the transition into the NMCI environment. The process of identification and prioritization has been largely delegated to the unit or organizational level. The customers, through their Claimant organizations, inform the DON CIO of what they need within NMCI to be fully functional. The DON CIO will then assimilate, along functional lines, those applications that are not obsolete, not insecure, and not redundant.

The intent is to produce horizontal integration among communities of interest (divided along functional lines) and the desired standardization and interoperability across claimants for each functional area. This transitional approach is an exhaustive bottom up process that is focused on claimants and their interactions with their subordinate organizations. While the intent is to achieve horizontal integration of claimants along functional lines, the

transition appears focused along traditional vertical organizational lines.

Placing the majority of the responsibility for identifying and prioritizing legacy applications at the unit level emphasizes the unit level functions. The hope is that if the unit level transition is done correctly, all the pieces will fall in place together when the transition is completed. If the effort is focused on the unit level, when do the intra unit functions get emphasized? The assumption is that they will be achieved in the aggregate as the entire enterprise comes together. This assumption may not be valid, however, and could result in gaps or incompatibility between Claimants. To be successful, this assumption relies on the idea that all units will identify all of the legacy applications in the same manner for prioritization for transition. This appears to be a very large assumption given the number of organizations being transitioned.

The vertical emphasis of the transition can be seen in how the PEO IT Office is tracking the NMCI transition. The figure below is from a briefing from the PEO IT office dated 12 April 2001. This information was presented as an update of the transition process and illustrates the vertical emphasis of the NMCI transition plan **[PEOIT01]**.

CLAIMANT SITES	Original AOR Date	Projected AOR date	Actual AOR Date	STATUS
NAVAIR				
NAS Pax River	15-Dec-00	12-Jan-01	12-Jan-01	COMPLETED
NAWCWD China Lake	15-Jan-01	16-Jan-01	16-Jan-01	COMPLETED
NAWCWD Point Mugu	15-Jan-01	16-Jan-01	16-Jan-01	COMPLETED
NAWC-TSD Orlando	15-Jan-01	16-Jan-01	16-Jan-01	COMPLETED
NAEC Lakehurst	1-Feb-01	1-Feb-01	1-Feb-01	COMPLETED
NAWCWD White Sands	1-Mar-01	1-Mar-01	1-Mar-01	COMPLETED
NATEC North Island	1-Mar-01	2-Apr-01	2-Apr-01	COMPLETED
RESFOR				
NAF Washington	15-Dec-00	3-Jan-01	3-Jan-01	COMPLETED
NARC Lemoore	15-Jan-01	2-Feb-01	5-Feb-01	COMPLETED
VFC13 Fallon	15-Jan-01	2-Feb-01	5-Feb-01	COMPLETED
NAS Atlanta	1-Mar-01	9-Mar-01	9-Mar-01	COMPLETED
REDCOM South HQ	1-Mar-01	21-Mar-01	23-Mar-01	COMPLETED
N&MCRC Dallas/FT Worth	1-Mar-01	21-Mar-01	23-Mar-01	COMPLETED
N&MCRC Waco TX	2-Apr-01	2-Apr-01	2-Apr-01	COMPLETED
N&MCRC Shreveport (Bossier city) LA	2-Apr-01	2-Apr-01	2-Apr-01	COMPLETED
N&MCRC Austin TX	3-Apr-01	3-Apr-01		ON TRACK
N&MCRC Little Rock AR	3-Apr-01	3-Apr-01		ON TRACK
N&MCRC San Antonio TX	5-Apr-01	5-Apr-01		ON TRACK
N&MCRC Tulsa (Broken Arrow) OK	5-Apr-01	5-Apr-01		ON TRACK
NRC Hartingen	6-Apr-01	6-Apr-01		ON TRACK
NRC Oklahoma City	6-Apr-01	6-Apr-01		ON TRACK

Figure 5 Current Transition Plan

The transition plan is organized along claimant lines of responsibility and their subordinate organizations when examining the problem of legacy application transition. The existing plan focuses on organizational structure and doesn't specifically address the linkages that exist between claimants in a primary way, only as a secondary event. The potential exists for the lateral interaction between claimants to be missed; resulting in remedial effort to produce the required functionality after transition is complete. Given the structure of the NMCI contract award, this would likely result in additional expense. To prevent such an occurrence an organizational method for the transition must be chosen that cuts across claimants.

The transition of legacy networks and applications should be approached in the same manner as the evaluation of the entire network system, from the view of mission function. From reviewing the steps included in the legacy transition guide and the memorandum from the DON CIO it appears the NMCI mission is not a primary consideration in the legacy transition plan. The unit level organizational format employed was responsible for bounding the legacy problem but it is unlikely this approach will help solve it. Therefore, another approach is required.

The mission function of NMCI is the variable that should determine which applications are transitioned or migrated to NMCI first. The selection of mission function is relevant because it is central to the concept of survivable network systems analysis method. As previously discussed, the NSA method is completely relevant because of the need to include legacy in NMCI during the transition and in the future. Handling the triage of the legacy applications with a mission oriented focus will cull the large number of applications of a basis that is relevant to the entire Navy and Marine enterprise, the mission of the NMCI network, and the war fighting mission of the Navy and Marine Corps. This is the most immediate problem to be handled by the PEO IT and the NMCI ISF.

The problems of security and interoperability are significant and relevant that will require considerable time and effort. They are in this case, however, secondary to the prioritization and focus of the transition efforts. If applications are transitioned in a random manner that does not consider the mission function of the network, the result is secure and interoperable applications on an individual or disaggregate level. The applications could be

safe and functional, but their contribution to the core missions of the Navy and Marine Corps team will be more obtuse. Interoperability and security must be dealt with for each application migrated into NMCI.

It is therefore more logical to spend the time, money and effort on transitioning and making secure those legacy applications that will serve your core mission functions first, and others that do not at a later time. Doing so will optimize the transition effort while building in the survivability characteristics needed to improve the system security required by the inclusion of such legacy applications. The same concept applies to the connection to external networks in the unbounded environment that NMCI will operate in. The effort should be on the network connections that serve the core mission functions of NMCI and the Navy and Marine Corps first, and the organizational functions second.

The more significant factor to the success of the entire enterprise network is the mission functionality provided by the transitioned legacy applications and networks and not the organization or organizations that possess them. Once the applications are selected for transition based on mission function, more efficient and effective methods for implementing the transition can be examined.

F. DESIGN METHODS FOR IMPLEMENTING LEGACY APPLICATION TRANSITION

The entire NMCI transition process can be viewed in terms of software design methodologies. There exist a large number of independent and interdependent requirements that must be organized and given cohesion to form the

complex, functioning system know as NMCI. Software design methodologies attempt to break the desired functions down into their simplest components and then, in a building block fashion, reassemble the functions to constitute a fully functional system of interoperable modules, each of which retains functionality that serves to produce the objective results (mission function) desired from the entire system.

Placing the NMCI mission function as the determinate for the entrance of legacy applications and networks we can now examine how, in terms of design methods for implementation, the current transition is designed and how the transition could be designed under the NSA method.

G. THE CURRENT TRANSITION; GRAND DESIGN METHOD

The current transition plan can be likened to the grand design method of software development. Under the grand design concept, requirements are captured up front and the total system designed in a single process. This process is illustrated in the figure below. The Grand Design method is a sequential method that relies heavily on the understanding of the system requirements at the inception of the project.




Figure 6 Grand Design Method

This method provides good results if from the start the requirements are well known and all processes fully understood. The difficulty with the grand design method occurs when there are functions or processes that are overlooked or under appreciated during the requirements development phase. The vertical emphasis of the current NMCI transition plan, organized along lines of claimants, does not emphasize the requirements for horizontal interactions between claimants and their subordinate organizations directly when examining the problem of legacy application transition. While meeting the need to connect all the Navy and Marine "nodes", the potential exists for remedial action by the contractor to provide the functionality missed in the transition phase of NMCI roll out. This requirement translates into lost capability and additional cost.

The implementation of the NMCI transition via the grand design method should meet the NMCI mission of connectivity. The unintended consequence of this method, however, is that the PEO IT office is left with an extremely large problem of how to separate the roughly 37,000 legacy applications [GCN02]. The application of the same grand design methodology to these legacy applications would require the NMCI ISF to determine how to transition all the applications in one very large process. While not being a single process, it is a series of repetitive efforts done based upon claimant and ignores the mission function the application serves within the network. This has proven to be impractical if even doable at all. The transition efforts may be focused on applications that have little or no consequence to the mission function of NMCI.

Implementing the legacy application transition along the same claimant lines is a possibility, but it would require claimants to be prioritized among one another by some metric. If an agreeable measure was at hand, transitioning a single claimant at a time ignores the horizontal integration that some of the legacy applications may possess. The true benefits of the legacy application transition would not be realized until all claimants that have a horizontal interaction with the legacy application have completed transition. This is at best inefficient. Certainly the benefits of the legacy transition will be wanted sooner rather than later.

When viewed in the context of the survivable network analysis (SNA), the grand design method ignores the core mission functions of the Navy and Marine Corps. Recall that for a network to be survivable it must continue to provide its primary mission functions even when suffering from an attack or a component failure. The grand design method does not support this requirement. The grand design method attempts to capture, design and implement all mission functions in a single effort. This implementation method assumes all missions are of equal importance, which is known not to be the case.

To meet the survivability requirement under the SNA method, the implementation of legacy applications needs focus on the mission prioritization of the network, just as in the application selection process. Implementing the transition to meet this demand can be enhanced by the employment of an iterative design method known as the spiral design method. When the spiral design method is coupled with the concept of survivable network analysis method the result will be a layered approach that build that

transitions the mission essential applications and networks first.

H. TRANSITION UNDER THE NSA CONCEPT; SPIRAL DESIGN METHOD

The spiral design method is meant to provide a series of iterative software developments for an application. The concept is used in the development of weapon systems software because of the complexity of the processes and the coupling that often exists between them. The advantage of this method is that the software functions can be designed in a logical series of iterations of the same process. Successive iterations of the design process increase the complexity and functionality in a layered manner. As each layer is completed it can be checked for errors and the functions validated more easily because of the smaller size of the code modules. Successive layers can then be completed and checked for function more accurately because of the understanding gained from the previous iteration of the software. It is essentially a building block approach to the software development process.

The application of the spiral development method to NMCI transition could provide the NMCI ISF and PEO IT a building block approach to the construction of the mission functions for NMCI. The spiral method creates a logical division of function and complexity so that the relevant warfare commanders in coordination with the PEO IT can select and prioritize the legacy applications to be transitioned in a manner that supports the primary mission of the network and the relevant core mission of the Navy and Marine Corps team. The logical division of functions can be viewed as the separations of mission essential and

non mission essential functions to be performed within NMCI.

The present transition is not oriented in a manner that truly supports the war fighting mission of the Navy and Marine Corps. To date the focus is on connectivity. In chapter two of this thesis the case was made for a mission focus for the NMCI transition to pursue, that of Force Projection. The suggestion in chapter two was that NMCI should focus first on the logistics mission and then upon readiness. Applying these mission areas as the focus to the spiral development of NMCI will build the functionality within NMIC consistent with the layering of mission essential functions under the NSA method. The net result of this could be as each spiral is completed, the total force gains mission functionality that is relevant the core mission functions and capabilities of the network that are aligned with a core mission function of the Navy and Marine Corps.




Figure 7 Spiral Design Method for NMCI Implementation

1. The First Spiral; Roll Out (Connectivity)

The starting point for any spiral application must be established in some way, either a date or a capability. The reference line or starting point for the spiral method of legacy transition can be viewed as the contract award date. The transition steps taken to date can be viewed as the first iteration of the spiral method. As the transitions progresses, more Navy and Marine organizations are brought into the NMCI environment.

To date approximately 15 percent of the Navy has been transitioned into NMCI [RM02]. Once complete, all the Navy organizations will have achieved connectivity to NMCI at some level. The problem now facing the PEO office is what to do with the extensive collection of legacy applications. The answer to this problem is to first determine the primary mission for NMCI and second to prioritize within that mission the most significant function. As previously discussed it is my belief that the Force Projection mission is the most relevant to NMCI. Within the Force Projection mission, logistics is the most significant function and should be transitioned first.

While only a limited percentage of the units have been transitioned to NMCI, it may still be possible to inject the mission emphasis into the transition. With 15 percent of the targeted units cut over, there is still 85 percent that have not. The difficulty at this point may be the organizational resistance to realigning the transition. It may, from an organizational view be better to wait until cutover is complete before this approach is applied but, the gains in efficiency that could be achieved may outweigh the changing of horses at this time in the transition. The time lost would be equal to only that required to define

their mission focus and then redirect the teams and their efforts for the transition of the legacy applications. Some organizations may be delayed in their date of transition but the mission functionality gained from the reorganization should outweigh the effort required. At the end of the day, should the transition not be refocused, it would not preclude the execution of the second spiral, that of implementing the logistical mission functions within NMCI.

2. The Second Spiral: Logistics

The transition of the logistical system legacy applications is driven by the pivotal influence of logistics on conflicts. Before September 11th 2001, the United States had not fought a conflict on home ground since the Civil War. In every other conflict since then we have had to support our forces from afar, necessitating an extensive supply network. While the war on terrorism is being waged both at home and abroad, the importance of the logistical mission cannot be understated. The Battles of the Pacific and Atlantic during World War II show how the logistical battle must be fought and won first, before combat forces can engage with the hope of prevailing. The logistical mission within the U.S. in support of the current war is no less important. Successful response to terrorist events within CONUS can be equally enhanced by a logistical mission focus within NMCI. This necessity of logistics for effective Force Projection is then obvious. The additional benefit to this is the coupling that exists between with logistical mission functionality and the capabilities of COTS applications prevalent in the commercial sector.

The logistical functions of the Navy and Marine Corps share the greatest commonality with the private sector and present the most likely candidates for transition to NMCI. The logistical mission function is in essence an extremely large inventory and delivery system. The supply system is comprised of a hierarchical collection of supply points, each possessing a similar stock and line item inventory dependent on where they lie in the hierarchy of the supply system or where they exist in the geographical terms to those units they support. The supply system also provides a delivery mechanism, either DoD or commercial, for all of its customers. For this they possess a shipping and delivery system to ensure the material that is ordered arrives to the customer and can be traced while enroute. Neither the inventory nor delivery system is unique to the Navy or the federal government.

There are several major commercial companies that are capable of performing either of these functions. Because these functions are prevalent in the private sector they offer tremendous opportunity for expedient transition or migration to newer software. They may also offer the greatest cost saving to be found in the implementation. The private sector software applications and organizational scheme can be emulated and could produce greater efficiency than is presently available within the supply chain. Transition of logistical legacy applications takes advantage of the commonality with contemporary business practices and the relative maturity of the field of software development for inventory management and control and delivery. Transition of the logistical legacy applications also fulfills a primary mission function of the NMCI mission of Force Projection, which directly

supports a core mission function of the Navy and Marine Corps team.

3. The 3rd Spiral: Readiness

The third iteration of the spiral design should be for the transition of the readiness relevant legacy applications. The most significant of the readiness applications is the Status of Readiness and Training System (SORTS). This database system holds readiness data for every active and reserve component organization. This application is not yet web enabled but could likely be. This program may even be easily migrated into new software given that it is essentially a simple data base application. The SORTS application is however just the tip of the readiness iceberg. There are tremendous amounts of information that flow from detachments, squadrons, ships, air wings, and entire battle groups on a daily basis to a large number of Navy organizations around the world. Aviation Maintenance Readiness Reports (AMRR), Daily Operations Summaries (OPSUM), Logistical Requests (LOGREQ), and Casualty Reports (CASREP) are just a few of the data messages that are required on a daily or routine basis to the many supporting commands spread around the globe.

Some of these messages contain both readiness and logistical data. All of these traverse the Defense Message System (DMS) and arrive at their destination with data that is time late, but time critical. The migration of this data into a web environment would allow the support organizations to view it in real time. The information could be hosted in onboard (primary) and ashore (backup) data bases and accessed by the reporting unit when updates are required. Hosting the data on shore would reduce the

requirement for transmission over the DMS for routine information and given CONUS or other users the ability to access the data easily through existing web methods. At a minimum, organizations within CONUS would no longer require a DMS transmission from deployed units. Deployed units would require only the bandwidth necessary for updates. Fleet and theater commanders slated to gain the participating units can view their readiness status and anticipate their needs or assess their ability to participate in operations on a near real time basis. The visibility of the readiness data throughout the deployment cycle gives commanders at all levels a greater ability to assess the capability of their units on a near real time basis. The availability of readiness data within NMCI again supports its primary mission of Force Projection and provides Navy and Marine Leadership with near real time data in support of a core mission function for the Navy and Marine Corps.

4. Spiral Sub-Flows

A further advantage of the spiral method over grand design is the granularity to which it can be employed. Larger processes can be decomposed into a series of smaller spirals or related steps. The logistical spiral may be more accurately portrayed as a series of smaller spirals focused on a particular hull type, type/model/series aircraft, or carrier battle group. For example, the aviation supply functions for the Fleet could be transitioned by aircraft type (F-18, F-14, H-60F/H/B) or functional organization (CVBG, ARG, CVAW, MAG). These smaller iterations could be viewed as sub-flows in the larger context of the transition of the Force Projection

mission legacy applications. While the sub-flows would be essentially vertical (within a single community, type/model/series aircraft, hull type, etc) they would possess more focus in application than the larger flows developed under the grand design method of transition employed today. The focus of the smaller sub-flows is the same, transition of the logistical mission function. Figure 3-5 illustrates the iterative cycles for the force projection mission area.




Figure 8 The Spiral Method View of Force Projection

By the same token, this method can be applied to NMCI in an aggregate manner as more missions are required of NMCI. Successive spirals could add mission functionality in the same layered manner. The advantage again being that those responsible for implementing the functionality can separate the implementation in a logical manner and plan the

progression of functions in a manner that supports the larger mission requirements of the Navy and Marine Corps.

I. CONCLUSION

To take NMCI to the level of supporting the war fighting capability of the Navy and Marine Corps there needs to be a more specific mission requirement. The function of the network must be tied to a core mission of the Navy and Marine Corps. As discussed under mission analysis in Chapter Two, the Force Projection mission best fits the capabilities of the NMCI architecture and the mission needs of the Navy and Marine Corps team. Defining a specific mission for NMCI will define what functionality it must possess to be an effective system for supporting the Navy and Marine Corps war fighting capability.

Defining mission function requirements for NMCI will also provide a method for separating the very large number of legacy applications into more logical, cohesive and manageable groups. This separation of legacy applications into mission functions also begins the construction of the layered mission essential functions required within NMCI. With the legacy applications prioritized for transition by mission function the most important applications are brought into NMCI first, supporting the core mission of the Navy and Marine Corps. The same method can then be applied to the existing applications as each layer is built through the spiral process.

The additional benefit of this method is that the security requirements necessary for the inclusion of legacy applications and networks are handled in the same logical manner. The net effect is that NMCI could begin to develop survivability characteristics that it would not otherwise

possess under the existing transition plan. At the end of the current process, if it continues as planned, NMCI will have internalized some of the legacy applications and networks without any additional effort to mitigate their presence.

Coupling the mission focus of the transition with a more efficient method of implementation (spiral design) will increase the efficiency with which legacy systems are handled and NMCI mission functionality is increased. Focusing transition on mission functions eliminates the requirement for remedial efforts at the transition after the initial effort has been completed. Moving a complete mission function in a single effort enhances the functions of both the network and the Navy and Marine Corps team.

For NMCI to be effective, it must include legacy applications that are essential to the present organizational functions of the Navy and Marine Corps support elements. Application of a mission focus to the transition which is implemented through a spiral design method will produce a more directed approach to culling these legacy applications for inclusion in NMCI. It will also induce the necessary survivability characteristics that will compensate for existence of the legacy applications and networks.

This method may not make the actual transition of applications easier on an individual basis, but it will reduce the complexity of the selection process, give the transition process greater focus, begin the growth of survivability characteristics within NMCI, and create the direct linkage between NMCI and the war fighting requirements of the Navy and Marine Corps team that is

necessary for NMCI to be successful both now and in the future.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. NMCI SECURITY ARCHITECTURE

A. INTRODUCTION

When examining the security architecture of any network there are two essential elements to be considered. The first is the strategy that is employed for the defense of the network and or the data that flows within it. The second is the tactics employed in the execution of that strategy. These two elements combine to produce the overall network defensive (security) architecture.

The strategy of the network defense is centered on what the ultimate goal of the efforts of defensive systems is to be. Surveying the field of enterprise network implementations it can be seen that network defensive strategies fall into one of three general categories: protection of the network nodes and links, protection of the network data, or protection of network access or availability.

A network defensive strategy centered on the protection of a network's individual members is focused on the continued operation of the individual hosts and servers. This strategy is prevalent in the majority of enterprise wide network implementations and shall be referred to as a "hard target" defense. Most of the enterprise-wide networks contain a wide variety in content of communications and possess as their primary function the connectivity of all the nodes within the enterprise. Under this strategy, the ability of individuals to communicate is considered to be of greater value to the enterprise than the content of each specific communication. By emphasizing

the protection of the individual hosts, it is hoped that the network will remain viable and usable to all. To support this strategy, organizations deploy layered levels of hardware and software in a distributed fashion across the network.

The intent is to give all network hosts an equal capability to resist attack. This includes intelligent agents that serve to alert the entire network should an attack materialize. The emphasis of the survival of the individual node under this network strategy is synonymous to the emphasis on the survival of the individual soldier applied in military strategies.

While not exclusively so, real-world military strategies are constructed around the idea that for the greater benefit of all, the defensive efforts should be focused on the preservation of the members through the execution of mutual support. Infantry units employ this strategy through the use of interlocking fire plans and preplanned artillery barrages. The survival of the members on both the front line and in the rear areas is dependent upon the successful coordination of all their efforts in the execution of a cohesive defense. This allows the entire organization to resist while maximizing the chances of survival for all the members in essentially equal fashion.

The retention of network data is a network defensive strategy that focuses less on the survival of the individual node and more on the information the nodes process or temporarily retain. Organizations that depend heavily on the validity, accuracy, and currency of data employ this strategy. Banks are a good example of such organizations. While seeking to avoid compromise that

would limit their ability to provide their service, they are more concerned with the integrity of their data. Their data is at the heart of their mission of commerce, and the loss of it could be fatal. For this reason these networks emphasize redundancy in data storage and extensive and reliable backup capability. The same data may be stored in multiple locations across the enterprise and or across the country. These organizations are willing to accept the loss of servers, hosts, and web site access in exchange for the preservation of the information that is essential for the ongoing function of the enterprise. Breaks in service, while undesirable, are often temporary and impart significantly less damage.

Protecting network access is a strategy that is often employed by on-line retailers or service providers. Their data functions are not time critical and their monetary functions are typically handled through third parties, so the critical element to the ongoing operation of the enterprise is the ability for customers to access the organization's virtual sales counter. On-line retailers and auction houses are examples of these types of organizations. Periods of loss of service or accessibility translate into lost revenue and directly impact the ability of the ongoing enterprise.

These organizations therefore emphasize availability of their web site through the use of redundant host locations and reliable crossover between them. If the enterprise web site hosted in a particular region is lost, an adjacent regional host will assume the load so that the entire customer base can be continuously served. This serves to minimize non-availability while maximizing the opportunity for the organization to obtain revenue. The

primary concern of the enterprise, access by and to customers, is thus met through this strategy.

The second element of network defense consists of the tactics used in the application of the network defensive strategy. The hardware and software systems deployed within the network and its resident hosts and servers constitute these tactics. These systems constitute what has been traditionally thought of as security. The importance of these two elements of network defense is that for the defense to be effective they must be aligned and consistent with one another. If not, then there are likely weaknesses within the network defense that could be exploited.

The state of the art in network defensive software and hardware are represented by intrusion detection systems (IDS), firewalls, anti-virus software, link encryption, and virtual private networks. These are the tactical systems that are deployed in a variety of ways in an attempt to ensure authenticity and confidentiality within a network.

Implicit with the evaluation of network security should be the examination of availability and quality of service. The implications of these factors (availability (Ao) and quality of service (QoS)) are frequently overlooked but carry great significance. These levels of availability and quality of service provided can influence a network's ability to resist attack and recover from its effects.

Availability within a network switched environment is defined as the accessibility of input and output ports. The significance of availability is that if not adequately provisioned, it can become the Achilles Heel of any network security strategy. Secure networks that possess inadequate availability can become isolated through exploitation of

single points of failure within the network structure that may not be adequately addressed by the network security tactics. Availability then is the examination of the network structure for multi-path routing to ensure reliable access to the input and output ports.

Quality of service is the ability of a network to provide better service to selected traffic that is flowing within the same network. This can be referred to as differential service. The primary goal of QoS is to provide priority including bandwidth, controlled jitter and latency, and improved loss characteristics to selected traffic during times of restricted bandwidth or availability [CISCO02]. This can be contrasted with the existing "best effort" service employed within the Internet and most other networks where all traffic is treated with the same priority for bandwidth, jitter, latency, and loss characteristics.

Networks that do not possess the ability to provide differential service can find themselves unable to communicate important information during periods of restricted bandwidth or availability. This is the result of the internal inability to discriminate among the high and low value information flowing within the network and then allocate the available bandwidth accordingly.

In summary, examining the security architecture of any system means that all of these factors should be considered, strategy, tactics, availability, and quality of service. Determining the strategy and tactics employed for network defense, coupled with the level of availability and quality of service, will determine how effective a network will likely be in the event of a compromise or intrusion. This is a much more useful determination than merely

examining traditional network security measures. The traditional definition of security extends only as far as the defensive mechanisms that constitute the network perimeters or define the enclaves. The strategy, tactics, availability and quality of service of a network combine to determine the survivability of a network. The Carnegie Mellon University Software Engineering Institute defines network survivability in similar terms [RJE98] and it is this frame work that will be used for the examination of the NMCI security architecture.

B. NMCI SECURITY STRATEGY

The NMCI Information Strike Force (NMCI ISF) has employed a defense-in-depth concept of layered security measures for the protection of NMCI. This layered defense is an enclave-based security strategy aimed at providing the desired level of information assurance by providing high resistance to attack while minimizing the security weaknesses of any particular security component within the defensive mechanism [RAY01]. This explicit definition of NMCI security strategy maps onto the strategy focused on the preservation of the individual nodes discussed earlier. Further examining the NMCI enclave approach to security it can be seen that this strategy is host- (seat-) centric, with the innermost layer of network defense at the host level. Figure 4-1 gives a representation of how the enclave strategy operates in coordination with the deployed boundary systems within NMCI.

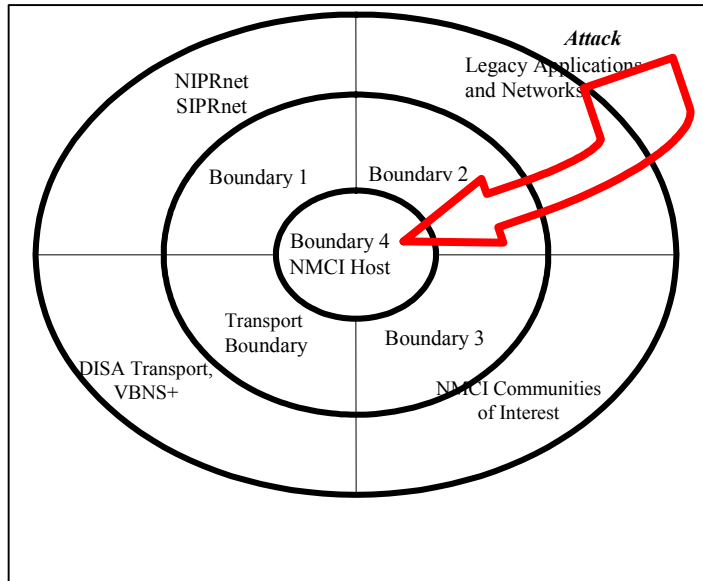


Figure 9 NMCI Seat-Centric Network Defense

The boundary system places the NMCI seat as the final layer of defense should an attack be attempted against the network. For an attack to be successful against NMCI, this strategy attempts to force an intruder to first penetrate the outer ring (legacy, NMCI Community of Interest, etc.) before being able to then assault the next boundary, and then subsequently the NMCI host or server. The intent of this defensive strategy is to force an attacker to fritter away his time and energy attempting to move from one area to the next through the layers of security, inward toward the host or server. At each level there are deployed hardware and software meant to confound the attack and warn the network of a potential or actual attack.

This strategy reflects the experience gained through the use of the hard target approach to network defense. This strategy is founded in a principle put forth by the great Prussian military strategist Carl von Clausewitz.

If you entrench yourself behind strong fortifications, you compel the enemy to seek a solution elsewhere.

Carl von Clausewitz, *On War*

The hard-target concept relies on an attacker to do one of two things when presented with a formidable network defense. The attacker will choose to attack another less well defended target, one that can be overcome within the attackers existing capabilities, and ignore the hardened target. The second option for the attacker is to assault the more formidable defenses of the hardened target and be depleted and or defeated in his attempt.

There are weaknesses in the hard-target strategy and its application to network defense that designers have failed to see. They have overlooked the unstated importance included in Clausewitz's strategic advice. If the enemy is compelled to seek another solution as Clausewitz suggests, then the defense must consider that solution and develop an effective response. The failure of the hard-target strategy within NMCI is its ability to effectively address the enemy's other solution in two ways.

First, the hard-target strategy cedes the initiative to the attacking force. While Clausewitz's statement recognizes the value of fortifications to the defender, he likely never advocated their use solely in defense. Defensive fortifications are only as valuable as the overall defensive strategy is able to flexibly respond. The Maginot Line and the fortifications at Eben Emel are perfect illustration of this shortcoming of over-reliance upon static defenses. The Maginot Line and the fortifications at Eben Emel were defensive solutions that

were rooted in the past. They relied upon history to tell them what form the next attack would take. The Germans, having learned the lessons of World War I, sought another solution. The German Blitzkrieg and airborne warfare were the new forms of attack. The French and Belgians relied too heavily upon their static defenses and failed to consider the potential German solutions. They sat and waited for the assault.

The forts at Eben Emel, while correctly placed and equipped, were unprepared for an airborne assault and so defeated in detail by German paratroopers and glidermen that attacked from the sky. The Maginot Line was even less relevant to the defense of France. Once flanked, the fortifications of the Maginot Line became irrelevant to the fight at hand as the German Armies rushed to the open fields of France. Having allowed the Germans the initiative, and not possessing a thoughtful plan to counter the German attack, the French defense became disintegrated and ineffective at the strategic level, the results of which are obvious to all. The implication for NMCI is the same.

The second failing of the hard-target strategy is that it does not address the mission functions of the network. Armies of Clausewitz's time and beyond have always possessed as their primary guidance some mission, however defined. Implicit with armies and not with networks is an inherent offensive capability. The designers of the hard target defense failed to recognize the significance of the offensive component of any army, even while executing a strategic defense. In military operations, the desire to execute offensive operations (attack) is in tension with the need to retain a defensive capability. This is not the

case in the networked environment. Network attackers do not have the burden of being required to react to a counter move by defending forces. There is little if any meaningful threat to them. They have the luxury of pursuing the attack until exhausted or effectively shunned, but they live to hack another day. The hard target defensive strategy gives NMCI nothing in the way of an offensive capability that would enhance its mission effectiveness in the face of any active intrusion or attempt. NMCI, like much of the French Army in the spring of 1940, sits and waits.

C. INITIATIVE CEDED TO THE ATTACKER

The traditional sense of network defensive strategy employed within NMCI concedes the initiative to the attacker. There are no proactive measures taken as part of the defense meant to deter an attack or enhance the mission performance of the network while under attack. Deterring an attacker requires that the defending party possess the ability to either deliver greater punishment to the attacker, or possess the ability to deny the attack or any hope of having meaningful effect on the target. There are legal issues that restrict the former behavior from occurring, so while desirable, they are impractical and illegal in the current environment. The latter behavior, blunting the effect of an attack, is not fully addressed by the capabilities of firewalls, intrusion detection systems, encryption, or virtual private networks that constitute the NMCI network boundaries. By examining these individual tools that constitute the bricks in the NMCI boundaries, their minimal deterrence and lack of mission enhancing capabilities within the network can be seen.

1. Intrusion Detection Systems

Intrusion detection systems are passive systems that examine network traffic and user behavior for a weakness that correlate to suspicious or malicious behavior. Typically, an intrusion detection system follows a two-step process. The first are host-based and considered the passive component, these include: inspection of the system's configuration files to detect inadvisable settings; inspection of the password files to detect inadvisable passwords; and inspection of other system areas to detect policy violations. The second procedures are network-based and are considered the active component: mechanisms are set in place to react to known methods of attack and to record system responses **[TT02]**.

The IDS strives to deter the attacker by eliminating known or recognized malicious behavior that is already documented as a security threat. The first step eliminates the known weaknesses and known behaviors, but successful attacks often are the result of newly developed malicious behavior and newly found weaknesses. Consequently, established malicious behavior is screened out, but assistance against new malicious behavior is very problematic. Independent IDS response will likely be limited to preprogrammed responses that are historically based. New malicious codes and behaviors often specifically attempt to avoid these pre-existing patterns to avoid detection. In either event, the IDS primarily provide the network with a notification of an intrusion and taxonomy of events in the aftermath. This response, while relevant, does nothing to assist the network in the performance of its mission while under attack, and in the

event entirely new tactics are employed against it, may not provide sufficient or any warning at all.

2. Firewalls

Firewalls are equally incapable of dealing with new behaviors. A firewall may be either a packet filter or a proxy server in nature and attempt to eliminate malicious behavior based on a set of rules that are implemented by a system administrator. These rules are essentially experienced based, making use of previous events in the determination of what type of behavior or packets can be viewed as malicious and should be shunned. What constitutes illicit behavior or an infectious packet is largely determined by previous events and this experience is what drives the development of the governing rules sets that the firewall uses to perform its job. The behavior of the firewall in the event of detection then is predictable, and so as long as the attacker avoids known behaviors he retains some potential for success.

In addition, the effectiveness of the firewall is limited to its ability to restrict access to the network, impacting users, friend or foe, equally. Restricting access within the network may not be desirable, or helpful in the event of an intrusion. The response is essentially experience-based and therefore reactive, doing little to assist the network against new malicious behaviors.

Also, the firewall's primary response to a threat, the restriction of access to ports, does nothing to improve the performance of the network mission. A firewall could resist attack by blocking all packet entry, but in doing so could provide the attacker with an effective "mission kill" of the network. Some ports must be left open, and the

implication is that they could be used by both friend and foe alike

3. Link Encryption

Link encryption, used extensively in the long haul movement of data in NMCI, is a mechanism to prevent traffic analysis, traffic flow analysis, or the jamming of transmission within a network. Recall the Buddenberg Matrix illustrated in chapter 1. Encryption makes invisible the sending and receiving node IP addresses and the content of the datagram. However, it only indirectly guarantees the confidentiality, authenticity, and doesn't even address the needs of the enclave defensive strategy. Over-reliance upon linked encryption for confidentiality makes the entire system only as strong as the weakest link in the encryption chain. While useful in overall view of security, they serve only to complicate the attacker's pre-assault efforts at detecting and localizing a target before determining the suitable method of attack.

In fact, depending upon the location of compromise, these systems could work against the defensive efforts to identify the source of the attack. The net effect is to make the attackers reconnaissance efforts more difficult, but once a successful attack has been executed linked encryption is lacking in its ability to mitigate the effect of an onslaught. In fact, continued use of a compromised encryption method is a counterproductive and inherently insecure act.

4. Virtual Private Networks (VPN's)

VPNs are a means of connecting geographically separated members of a network to the local area network in

a secure and confidential manner. Within NMCI, VPNs are used to connect members of communities of interest within NMCI, and to connect remote users to NMCI proper via a secure connection. The benefit of this application to the remote user is the ability to touch NMCI through the unwashed environment of the Internet from any desired location. The benefit to the network is to provide some assurance of authenticity and confidentiality within the network while providing a required service. It also provides some mechanism of security against traffic analysis or the jamming of transmission within a network traffic flow, in the same manner as link encryption.

However, traffic flow analysis may not be prevented completely as the headers of the IP data grams in VPNs are unencrypted, allowing an observer to at least identify the gateways used for the transmission. The utilization of the VPNs to support remote users also undermines the enclave strategy of NMCI by placing the host outside of the other boundaries. While operating remotely, the VPN capable host does not enjoy any support from the other NMCI boundary applications. VPN host defense capabilities are limited to boundary layer 4 and should those be compromised the other layers can then be bypassed through the illegitimate use of the connectivity provided by the VPN.

This act is known as "island hopping". Again, should a compromise occur, the utilization of a VPN offers no assistance in the continued function of the network. The VPN has no value added capability to help mitigate the attackers effectiveness, and if exploited could actually be assisting the attacker in his endeavor.

5. Anti-Virus Software

These applications function in a manner similar to that of firewalls. They possess the ability to identify and shun IP packets that possess characteristics that are known or suspected of possessing malicious content, or promote malicious behavior. These systems are widely deployed at the host and server level in an effort to screen out the transmission of harmful software applications. They, like firewalls are experienced based applications that require a historical profile of files that possess hostile characteristics. They are effective at eliminating items that have existing virus qualities, but their effectiveness fades quickly when presented with new software viruses or worms. Successful worms and viruses often masquerade as legitimate applications, or come packaged in a binary manner so as to avoid detection. If the inherent characteristics are not known or identifiable as hostile, they will be ignored. Since the application is only as smart as the most recent successful exploit, its usefulness in the immediate post attack environment is limited to preventing known virus or worms from exploiting the ongoing attack. Anti-virus software is reactive, and when faced with newly constructed malicious code often fails, providing no assistance to the network. In addition, the dexterity of the anti-virus software response and the policies that guide their actions impacts the usefulness of the network.

At present, Department of Defense DoD restrict executable as a virus threat. Consequently java applets and cookies are eliminated, as well as some spread sheet applications that contain "macro" level functions. These types of executable code are useful to the network members,

but they also fit the profile of potential virus hazards. The result is that both good and bad code are preclude from the network, reducing the overall usefulness of the network to its members.

6. Initiative Ceded to the Attacker: Summary

In summarizing the limitations of these tools it is obvious that they provide only a limited, reactive response to network intrusions that have been previously experienced or observed in some manner. In addition, they do little or nothing to assist the NMCI network in performing its mission in the face of the attack. It should also be noted that these applications are dependent upon the human intervention or other software applications to remain current and viable.

The overall effectiveness of the strategy is limited by the weakest link within the defensive chain. The individual applications are only as good as the management process that they each rely upon to keep them current when compared to the threat. Their effectiveness is a dependent variable based on the timeliness and accuracy of the updates initiated by a separate agent of the network. Second, the effectiveness of the strategy is limited by the ability of the weakest individual application within the defensive chain. The capability of the firewalls and IDS are negated by the weaknesses in the authentication scheme or the VPNs or the data base of the anti-virus software.

The net effect of reliance upon these applications for defense is yet again to cede the initiative to the attacker. The hard-target strategy within NMCI relies upon the attacker to stumble over one of the application trip-wires before any action is taken. This allows the attacker

to study the networks defensive systems and effectively plan his attack based upon their known, relatively static qualities. An unsuccessful attempt will illuminate the attacker, but he will suffer little else. The attacker gains from the failed attempt by evaluating the response and the defenses. NMCI likely gains little other than a validation of the prepared defense. The benefits to this trial and error approach lie overwhelmingly with the attacker. He is allowed to probe until he penetrates the defense, whereupon he gains tremendous knowledge, while the defender may be unaware.

This strategy and these tools may be effective at excluding a majority of potential attackers, those that attempt to exploit defined weaknesses employing existing tools, but they do not adequately address the need to counter those that do not follow the previously discovered path. The hard target defensive strategy grants the opponent the luxury of picking the time and place of attack, based on the knowledge gained through the examination of the relatively static defense, while the NMCI must rely on historical data to tell it about an opponent as yet unseen, or to protect a weakness that has yet to be discovered.

D. FAILURE TO ADDRESS THE MISSION OF THE NETWORK

Attempts to compromise networks frequently begin with an attempt to compromise a host, either internal or external to the targeted network, through an exploit or via a method of social engineering. Once compromised, the attacker makes use of the legitimate network resources provided by the host in an illegitimate manner to further

compromise the network. The advantage to the attacker is that by controlling a network host, they can bypass some of the network defensive systems. The result for NMCI could be that the integrated defensive system of boundaries would be decomposed into single barriers that are attacked in a piecemeal fashion. The hard-target strategy designed to defend the network by protecting the individual hosts becomes inverted. The network becomes only as well defended as the individual hosts are able to defend themselves. If the host level defenses fail, NMCI's hard target strategy doesn't address what to do next. This shortfall is the result of the enclave or hard target strategy's failure to consider the importance of the mission performed by the network.

As discussed earlier, the tools used in the traditional practice of network defense do not address the mission performance of the network. These defensive tools are responses to specific threats that have, over time, been improved to counter the increased capability of the attacker. Firewalls and IDS are the countermeasures deployed around the network that hope to screen out the offending behavior. As the attacks have gotten better, so have the countermeasures, and vice versa. The logic behind continuing to pursue this spiral is that the defenses will eventually outpace the improved attack capabilities, producing a condition of diminishing returns for the attacker. Unfortunately historical data tells us otherwise.

If the defensive strategy was improving at a faster pace than that of the attackers, then logically the number of compromises should be decreasing. Table 10 below shows that the number of cyber attacks has not ebbed as a result of our best defensive efforts. These numbers call into the

question the effectiveness of the purely defensive strategy in meeting the needs of a mission critical system like NMCI.




Figure 10 Reported Cyber Incidents 1990 to 2001

The problem lies in the misalignment between the mission of a network and the objective of enclave based network defense. Networks were designed to move data, not to defend it or protect it within a bastion or behind a rampart. The elemental design of networks is centered on making the movement of data increasingly more efficient. The enclave security strategy attempts to achieve the goal of security by employing tools that function in opposition to the elemental nature and primary mission of networks, movement of data. Attackers are employing a denial form of strategy when assaulting a network. The hard target defense within NMCI attempts to employ another form of the same denial strategy to defeat the attackers. To better illustrate this mismatch consider the application of Admiral Alfred Thayer Mahan's sea power strategy to the function of modern networks.

E. MAHANIAN STRATEGY AND THE NETWORKED ENVIRONMENT

Admiral Alfred Thayer Mahan was a noted naval theorist and considered to be the father of modern naval strategy as practiced by the United States Navy for over a century. Mahan's theory divides naval strategy into two bodies of thought. Those two communities are sea denial and sea control. Those that desire to deny use of the world oceans by others practice a strategy of sea denial. The sea denial strategy is intended to disrupt and deny the free use or travel of the seas by another. Sea denial is not considered a war winning or offensive strategy. This strategy is employed to try and mitigate the chance of losing the war. Those wishing ensure freedom of the seas or to dominate it for their use in furtherance of national aims practice a strategy of sea control. Those that practice sea control seek to dominate the sea, allowing them free rein to transit for commerce or the ability to project power on a distant land mass. Sea control is an offensively oriented, war winning strategy. Using these definitions, Mahan's concepts can be mapped onto network strategies and the mission of networks.

The Internet, or the networked environment, is the neutral medium that is used for the movement of data between the nodes. The network is the ocean upon which the data is moved from node to node or port to port. Those that wish to block the use of networks or the Internet practice an *information denial strategy*. They seek to disrupt or deny the free, unrestricted movement of information. Those seeking to get their data across the network in spite of the efforts of those executing an information denial strategy are following an *information control strategy* [RB97]. The enclave defensive strategy

within NMCI attempts to counter the attacker's information denial strategy by executing a similar information denial strategy of its own. This denying of the denial strategist is a purely defensive and avoidance strategy that assumes from the onset that the war can not be won.

This strategy is comparable that pursued by the Imperial Japanese during World War II after mid-1942. The Japanese had fortified many Pacific islands, hoping to deplete or deny the American attack by extracting such a high price in lives and material that they would in some manner prevail and retain their territorial gains in the Pacific. However, without the sea power to control the oceans in and around these island fortresses, the Japanese could not influence the movement of the American forces or pursue a war-winning strategy. Subsequently, American forces attacked at the points of their choosing and on their timetable, crushing some very formidably defended islands while selectively avoiding and ignoring others.

The lesson for traditional network defenders is that the enclave strategy is rooted in the same idea held by the defeated Imperial Japanese. The enclave defensive strategy creates an island chain of fortresses within the ocean of NMCI, believing that dominance of the internal network environment will manifest control in the external environment. The tactics of the enclave strategy do not enhance the connectivity or availability between the nodes. The enclave strategy only assists the network in resisting the attempts of others in the execution of an information denial strategy against it. Simply preventing the execution of an information denial strategy by an opposing force, however, does not equate to the application of an information control strategy for the network. For a

defensive strategy to be effective it must possess an offensive component. This offensive component must assist the network in the performance of its primary mission, moving data from node to node, in spite of any intrusion. Being able to do so effectively blunts the attacker's efforts and could provide some deterrent capability.

F. NMCI STRATEGY SUMMARY

What is needed, but is not present within NMCI, is an information control strategy that that supports the essential mission functions of the network and makes use of the enclave strategy defensive tools. The network strategy must assume a more offensively-oriented posture. This idea is a distinct departure from the traditional notion of network defense. An offensive network strategy must not be confused with the idea of computer network attack (CNA). CNA is an execution of the denial strategy against another network.

In this context, an offensive network strategy is one that employs mechanisms that will enhance the capability of a network to move data from node to node in spite of compromise or damage. The reliance on signature-based recognition applications is being rapidly overcome by the ability of attackers. Signature based recognition may become impossible in the future because of the innate ability of these malicious codes to morph or change [5]. Purely defensive tactics may have been overcome by the more capable "blitzkrieg" posed by these viruses and worms.

Intrusion systems may suffer the same fate. They rely on the identification of malicious behavior based upon historical profiles. Unfortunately, malicious behavior can result from what appears to be authorized, legitimate

behavior within a network. Rules and permissions are only as effective as the knowledge base of those that develop them, and they are only as good as the experience of the developer. The enclave strategy is fighting a losing battle, hoping that the information denial forces never become capable enough to totally dismember its defensive mechanisms. There needs to be an offensive strategy for NMCI, one that enhances connectivity, availability, and can distribute bandwidth in response to an attack, supporting the networks mission of data transfer. Employing a strategy that does anything less will leave NMCI wanting when faced with new capabilities of the information denial forces.

G. NMCI SECURITY TACTICS

The second element of any network that should be examined is the employment of the defensive applications or mechanisms to defend the network. The importance of this is that whatever the particular strategy is that has been chosen to defend the network, at the very least the defensive mechanisms should be arrayed in a manner that is consistent and supportive of that strategy. If not, an effective strategy may be able to produce only limited mission success and an ineffective strategy may become that much worse. In either case, the result is undesirable and regardless of the viewed viability of the network defensive strategy, the applications and mechanisms that constitute the tactical elements of the strategy should be positioned to optimize their performance within the network in support of that strategy.

As discussed earlier, any examination of network defensive tactics must include the evaluation of all three

of the relevant elements that constitute that defense. Namely they are availability, security, and quality of service. These three elements are interrelated and the shortcomings or failure of one impact the others, and consequently the network as a whole. For a defensive strategy to be successful, all three of these elements must be in alignment and adequately addressed so as to produce the desired levels of authenticity, confidentiality, security, and connectivity.

1. NMCI Availability

Within a network-switched environment, availability is defined as access to input and output ports. Availability goes directly to connectivity between nodes, which to date is the only clearly defined mission possessed by NMCI. For networks in operation there are definitive statistics. Availability is a very tangible parameter and there are web sites in existence that can provide very specific data concerning this parameter for hundreds of large distributed internet service providers around the globe. NMCI, however, does not yet physically exist in its entirety and so if any statistical performance on NMCI was available it would require some sort of rationalization to make a valid comparison. Therefore another method for comparing NMCI availability must be found.

NMCI is a services contract, and while there is no existing network that can provide historical performance data, a comparison of the performance parameters within the NMCI contract to the performance of other large networks is useful. An assumption can be made that NMCI contractual requirements equate to actual nominal performance of the network in its operating environment. Based on this

assumption, a valid comparison between NMCI contract performance and existing terrestrial networks can be made. The results of this comparison can then be used to evaluate the potential effectiveness of NMCI in the performance of its mission.

The nominal performance levels of NMCI can be found within the NMCI contract award. This services contract has thirty-nine specific service level agreements (SLA's) within it that define the desired performance characteristics of NMCI. Of these thirty-nine SLA's, twenty-six of them address availability in some manner. Of the twenty-six that address availability, ten of them are relevant in some manner to the operation of the network mission function of connectivity. Of these ten, there are eight SLAs that can be used in the comparison process in determining the adequacy of the availability provisioned within NMCI. These eight SLAs are identified by number and their contractual requirements summarized (in terms of availability) in the table below.

NMCI Service Level Agreement	Contracted Level of Availability
SLA 6: Web Access Services	99.5%
SLA 10: NMCI Intranet Performance	99.8%
SLA 11: NIPRNET Access	99.5%
SLA 12: Internet Access	98.0%
SLA 24: WAN Network Connectivity	99.99%
SLA 25: BAN/LAN Communication Services	99.9%
SLA 27: External Networks	99.5%
SLA 35: Information Assurance Operational Services (SIPRnet)	98.0%

Table 10 SLA Availability

However, to see the true value of the connectivity that has been provided we need to consider the totality of the connectivity is a function of the interdependent networks that constitute the enterprise network known as NMCI.

A complex intranet like NMCI is composed of a series of smaller networks. Under NMCI each of these separate internal networks has been provisioned individually, independent of the one another. The availability of the enterprise wide network to the user level then is a function of the availability of the series of independent networks based upon the distance from the host. Consider the following illustration.

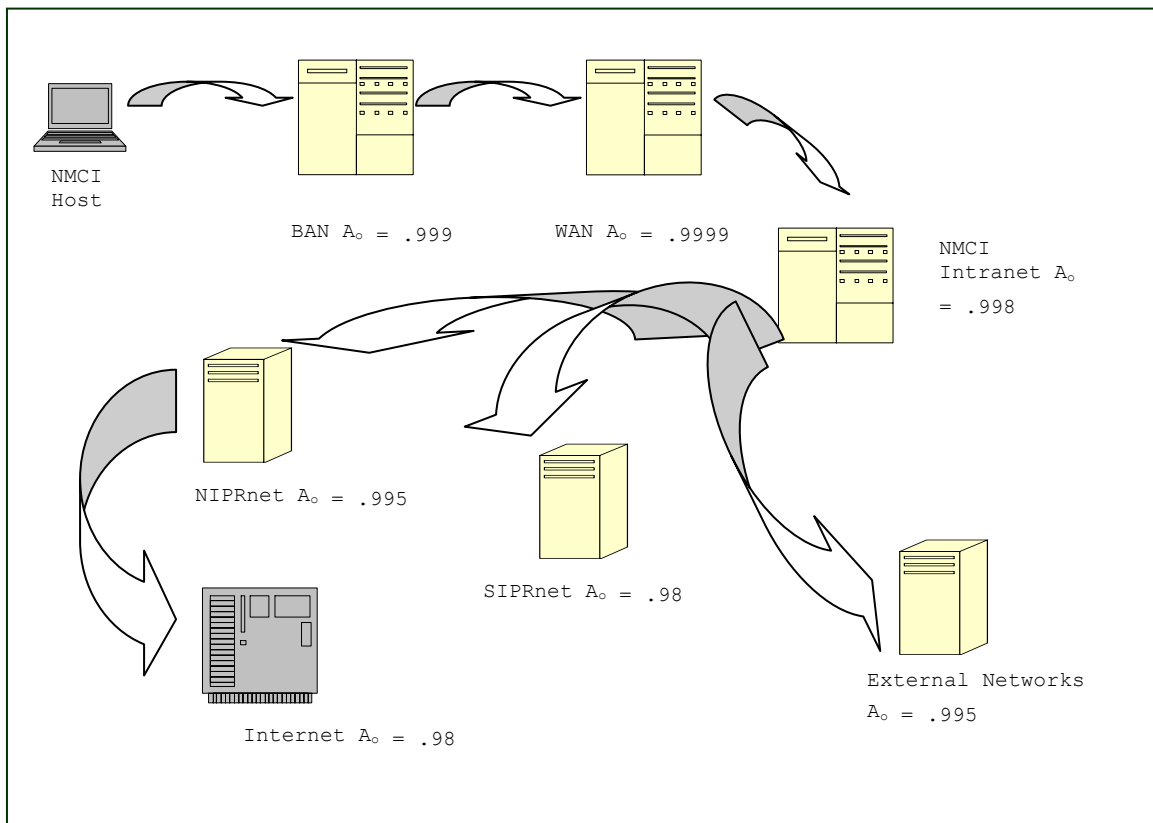


Figure 11 NMCI Series Network Effective Availability

The value that needs to be determined for a valid comparison is the effective availability at each level of

service under the SLAs when viewed from the user level. For this determination it will be assumed that host availability is essentially 100%, based on the examination being the accessibility of the higher level services. The effective availability of a host to reach each level is determined as follows;

$$A_o \text{ of NMCI BAN} = (1.0) * (.999) = .999$$

$$A_o \text{ of NMCI WAN} = (1.0) * (.999) * (.9999) = .9989$$

$$A_o \text{ of NMCI Intranet} = (1.0) * (.999) * (.9999) * (.998) = .9969$$

$$A_o \text{ NIPRnet} = (1.0) * (.999) * (.9999) * (.998) * (.995) = .9919$$

$$A_o \text{ SIPRnet} = (1.0) * (.999) * (.9999) * (.998) * (.98) = .97696$$

$$A_o \text{ of External Networks} = (1.0) * (.999) * (.9999) * (.998) * (.98) = .9919$$

$$A_o \text{ of Internet} = (1.0) * (.999) * (.9999) * (.998) * (.995) * (.98) = .97207$$

The calculation is viewed from the host level, based on the logic that the host level is where the vast majority of traffic will originate and end. This is also where the majority of the users will reside. It must also be noted that while the levels provided under the SLAs are higher, they are meaningless without considering the interaction required between the network segments within NMCI for the operation across the entire enterprise. Even though the WAN is provisioned at 99.99%, that fact is irrelevant when viewed in isolation because the WAN is of no value unless attached the users who reside at the BAN/LAN level internal and external to NMCI. The implication of this structure is that while SLA-provisioned availability is defined in the contract, the maximum attainable level of availability to the user is always going to be something less because the availability at the user level must take into consideration the down time accumulated by the cooperating networks segments. Table 11 summarizes the SLAs, their contracted

availability and their calculated effective availability when viewed from the user level.

NMCI Service Level Agreement	Contracted Availability	Effective Availability
SLA 6: Web Access Services	99.5%	97.207%
SLA 10: NMCI Intranet Performance	99.8%	99.69%
SLA 11: NIPRNET Access	99.5%	99.19%
SLA 12: Internet Access	98.0%	97.207%
SLA 24: WAN Network Connectivity	99.99%	99.89%
SLA 25: BAN/LAN Communication Services	99.9%	99.9%
SLA 27: External Networks	99.5%	99.19%
SLA35: Information Assurance Operational Services (SIPRnet)	98.0%	97.696%

Table 11 SLA Effective Availability

These effective availability numbers are valid for comparison because they reflect the down-time of the interdependent networks just as the data measured on the large internet service providers does. With these more representative figures we can make a comparison of the availability provisioned within NMCI and that provided by other, similar, large area networks and consider the implications this may or may not have on the performance of the NMCI network and mission.

A web site available hosted by the Matrix Netsystems collects data on the performance of internet service provides around the globe. This data is available for a variety of periods, but for the purpose of this examination the data for the periodic data used will be for the last 30 day period. This measurement approximates the period to be used for the evaluation of performance under the NMCI contract. The measurement performed by the Matrix Netsystems web site is done externally to the evaluated

network, representing the end users point of view. As such it considers any subdivisions that may exist within the evaluated ISP's and is representative of the performance of NMCI and its availability.

The Matrix Netsystems site subdivides the ISP evaluated by their size. For the sake of comparison, only the large ISP's that service within the United States were chosen for representative data. This was done based upon the similarity in size and the assumed similar U.S. terrestrial environment that NMCI will operate in. Given this as the starting point, data was obtained in May of 2002 for the twenty-six large U.S. ISPs measured by this site and a cumulative average calculated for the three parameters (latency, packet loss, and availability) measured. Based on that data, these are the averages to be used for comparison:

Latency 67.12 milliseconds

Packet Loss 0.2705 %

Availability 99.86%

Using this data the availability table can be reconstructed and the effective availability of the NMCI structure can be compared to that of a similar, large commercial ISPs.

NMCI Service Level Agreement	Effective Availability	Down time in Hours per year (A)	Commercial ISP Availability	Down time in Hours per year (B)	Increased (Decreased) Down Time in Hours per Year (A-B)
SLA 6: Web Access Services	97.207%	243.996	* 99.86%	12.23	231.766
SLA 10: NMCI Intranet Performance	99.69%	27.0816	* 99.86%	12.23	14.8516
SLA 11: NIPRNET Access	99.19%	70.76	* 99.86%	12.23	58.53
SLA 12: Internet Access	97.207%	243.996	* 99.86%	12.23	231.766
SLA 24: WAN Network Connectivity	* 99.89%	9.06	99.86%	12.23	(3.17)
SLA 25: BAN/LAN Communication Services	* 99.9%	8.736	99.86%	12.23	(3.494)
SLA 27: External Networks	99.19%	70.76	* 99.86%	12.23	58.53
SLA35: Information Assurance Operational Services (SIPRnet)	97.696%	201.277	* 99.86%	12.23	189.047

Table 12 SLA Effective Availability vs. Large U.S. ISP

When the two systems are compared on the basis of effective availability rates they appear at first to be only marginally different. The NMCI contract award is only marginally better in performance when compared to the existing ISP in two areas, specifically BAN and WAN availability. For these two SLAs the effective availability exceeds that of a commercial service. For the remaining SLAs the commercial ISP availability exceeds that provisioned by the NMCI structure at the contracted availability rates. The differences in rate, however, hides the more meaningful number, the total time the system is unavailable for use by the users.

To determine the total down-time accumulated for the system in a twelve-month period we need to do some simple calculations;

$(7 \text{ days} \times 24 \text{ hours} \times 52 \text{ weeks}) = 8736 \text{ total hours available in a year}$
 $8736 \text{ hours} \times (1 - \text{effective availability \%}) = \text{total down time for the}$
system

Looking at total down-time, there are obvious advantages to the level of service provided by the commercial ISP. Based on the assumption that the 99.86% availability is across the full spectrum of service, the ISP performs considerably better. Over a twelve-month period the commercial ISP accumulates less than 50% of the down time when compared to NMCI. This holds true for all of the services except the BAN and LAN. The BAN and LAN level service, however, is not reflective of the larger purpose of NMCI. NMCI is an intranet, meant to connect the entire Navy and Marine enterprise. The telling figure is when we examine how well NMCI can connect all of the members and supporting network functions.

The core of NMCI is reflected in SLAs 10, 11, 27, and 35. These SLAs address the capability of the system to connect all of the user members of the Navy and Marine Corps and the cooperating networks. When comparing the availability here we can see that NMCI is considerably less capable than the commercial ISPs. NMCI Intranet performance (SLA 10) accumulates more than twice the down time of the commercial ISP. The availability to the NIPRnet, SIPRnet, and other external networks (MCEN, IT-21) falls off considerably. The NIPRnet and SIPRnet expect nearly six times more down time, and the external networks more than sixteen times more non-availability compared to the service of a commercial ISP. This is surprising based upon the how NMCI was developed. In the process of

outlining the requirements for provisioning the network, the Navy consulted with large U.S. corporations, drawing on their experience of operating enterprise networks in the formulation of the NMCI contract [RM02]. If commercial ISPs can provide this level of availability, why doesn't NMCI meet that same level of performance? There are two possible answers.

The first possible answer to this question is that it is purely a business deal, unconnected to Fleet needs. The availability levels contained in the contract award were negotiated as part of the NMCI contract. The contractor may actually be able to provide service better than the negotiated availability, equal to the commercial ISP, and so the lower performance levels were negotiated to give them a margin for error and a potential greater profit based upon the performance incentives contained within the NMCI contract [PEOIT00].

The second possible answer is that because of fiscal constraints or other reasons this was the best that could be obtained under the contract. This answer, however, doesn't square with the pre-contractual discussions between the Navy and large U.S. corporations regarding enterprise networks and the provisioning of the NMCI contract [RM02]. It is illogical that the Navy would accept performance less than what is commercially available. Even when viewed in the best light, the Navy has contracted for availability within NMCI that is no better than commercial ISP service. If this is the case, the next question is whether that level of availability is adequate to support the mission requirements for the network.

To determine if the provisioned level of availability is adequate to support the NMCI mission, there needs to be

a clear definition of that mission. As was discussed in chapter 2 of this thesis, NMCI lacks a clear mission definition, but two things emerge from the examination of the mission statements that are made in the available documentation. First, NMCI seeks to achieve connectivity within the Navy and Marine Corps enterprise. Second, NMCI is determined to be a mission critical asset for the Navy as part of the larger global information grid used by the service. Using these two mission objectives was can compare the availability rates to determine NMCI's adequacy in reaching its objective goal of mission success.

It must be remembered that the NMCI contract is for services, the Navy and Marines have not dictated to the primary contractor any specifications of how to provide the service, only the level of service to be provided. As a consequence, then, if the contractor builds the physical network to these levels, NMCI will likely possess less than a dual threaded level of reliability. The implication of this is that there may not be reliable crossover and or adequate redundancy built into the network for the true needs. At present, the only level of the NMCI network provisioned to a dual-threaded level of availability is the WAN (99.99%). When viewed from the user level, the effective availability to the WAN is essentially equivalent to three nines (99.89%), but drops off when reaching the NMCI Intranet level to only two nines (99.69%). At this level, service is no better than single threaded and implies that the network possesses physical or logical single points of failure within its structure. Provisioning only single thread connectivity at the NMCI Intranet level gives only tenuous support to the stated mission objective of force wide connectivity. When the

need to interoperate with external networks is included in evaluating the adequacy of NMCI availability, things are no better.

As a part of the global information grid (the second stated mission), NMCI will need to be able to operate in a cooperative manner with these and potentially other external networks. SLAs 11, 27, and 35 address NMCI availability to the NIPRnet, SIPRnet, MCEN, and other external networks. The effective availability of these networks to the end user within NMCI is equally poor. Availability to the NIPRnet and the other external networks is well below the dual threaded threshold of four nines (99.99%), coming in at only 99.19%. Availability to the SIPRnet, the Navy and Marine Corps classified network system, is a dismal 97.696%. If the other primary mission function of NMCI is to be part of a global information grid, it appears that the information contained within that grid is of little importance to the Navy based upon the accepted level of availability to those networks. The relatively low level of availability to these other networks puts the usefulness of NMCI participation within the global information grid in doubt. The levels of availability provisioned for NMCI do not rise to the level of dual threaded capability. Instead, the capability is significantly less than that obtainable by applying relatively simple hardware redundancy to achieve dual threaded availability through multiple independent paths and reliable crossover.

a. NMCI Availability Summary

The effective availability numbers reflected in NMCI place doubt on its ability to achieve even the diffuse mission goals of connectivity and participation in a global grid in support of DoD. This situation is unlike the capability typically designed into any of the Navy and Marine Corps weapons systems today. Surface ships, submarines, and aircraft all some possess redundant capability within their basic structure. While not directly related to the combat capability of the larger system, these redundant features are meant to keep the ship afloat and moving, if for no other reason than self preservation of the members. Many military combat systems possess a capability that allows them to avoid becoming both combat ineffective and a combat liability simultaneously. NMCI does share this capability.

While not actually participating in combat, NMCI is a system meant to support the functions of getting the forces to the fight. The comparatively low effective availability within the network indicates that NMCI is at best a large DoD contracted ISP. NMCI is a system that could suffer damage or disruptions that would render it combat ineffective or a combat casualty as the result of an intrusion or compromise. Loss of NMCI could place vital assets out of reach of the command and control structure or render them useless because of the unavailability of data. The overall combat effectiveness of the Navy and Marine Corps team could be compromised because of the lack of system availability.

This means that NMCI will likely be ineffective at completing its intended mission of enterprise connectivity or participation in any global information

grid in anything other than a comparatively benign environment.

2. NMCI Security

In examining the security characteristics of a network, there first needs to be a definition of those characteristics. Once the desired characteristics are defined, then the network can be examined for their presence or absence, and a judgment made on the level of security that exists within that network. The characteristics that define security within any given network are;

- *Confidentiality.* Unintended recipients can't read traffic. Confidentiality includes secrecy of the data.
- *Authenticity.* Unintended originators can not fake traffic or forge messages. Authenticity is a superset of integrity.
- *Integrity.* Traffic hasn't been tampered with.
- *Non-repudiation.* Transmitted messages contain characteristics of attribution so that it can not later be denied.
- *Access control.* Unauthorized users denied use of network and computing resources.
- *Assurance of service.* The network is available for use and possesses resistance to denial of service attacks.
- *Traffic analysis.* Ability to derive intelligence from the addresses of messages, even if the contents are confidentiality-protected.
- *Traffic flow analysis.* Intelligence inferences gained by observing flows to and from commands and individuals.

- *Interceptability*. Ability of unintended recipients to receive traffic (regardless of whether they can read it).
- *Jammability*. Vulnerability of a link to interruption by signal interference. **[RB02]**

There are numerous ways of achieving these characteristics within a network and they can be applied to a multiple of layers within the OSI model. This is drawn out in ISO 7498-2 that lists the potential areas of application of security measures at each of the seven layers within the OSI model. The table that summarizes ISO 7498-2 is contained in chapter 1 of this thesis. The significance of this OSI model is that the actual implementation of the security measures must be resolved with the specific characteristics they are trying to impart to the network. The resolution process is summarized in what I earlier defined as the Buddenberg Matrix. The Buddenberg matrix aligns the technical solutions for security with their respective requirements and objectives. The great advantage of this matrix is that it permits an effective high level examination of the structure and software without requiring an in depth examination of the individual applications themselves.

ISO RM Layer	Requirements	Solution	Objective	Examples
7	Confidentiality Authenticity	Object Level Security	Object Level Security	S/Mime, secure shell, secure socket layer , VPN
3,4	Perimeter Protection of Enclave (Prevents DDoS attack)	Firewalls Intrusion Detection MAC/DAC	Secure the Network/Box (not the data)	Firewalls IDS Passwords
1,2	Traffic Analysis Traffic Flow Analysis, Jammability Detectability	Link Crypto LDI/LPD Spread Spectrum	Secure the Network pipe (transport)	KG-84 STU-III Wireless LAN

Table 13 Buddenberg Matrix of Security Requirements

The operative theory behind the Buddenberg Matrix is that in order to achieve the most efficient and effective security, all the elements (problem, solution, objective, and application) must be in alignment. This is not to say that applications can not be employed in other ways, but that to do so will at best sub-optimize the security for the network. Misalignment of requirements and applications within the Buddenberg Matrix produces inefficiencies and security that is likely less than thought or desired. Using the Buddenberg Matrix as an overlay to NMCI security we can begin to see there are potential gaps.

The NMCI security architecture is an enclave-based defense-in-depth concept that would employ services at OSI layers three and four. For the obvious reasons the network also utilizes applications at OSI layers one and two to provide resistance to attempts to jam or intercept the transmissions within the network. The boundaries that

constitute the layers of network security within NMCI rely heavily on many of the contemporary hardware and software applications for a security solution. These layers and their respective tools are summarized below.

Boundary Layer	Solution	OSI Layer	Requirement
Transport Boundary (Wide area Transport)	Link Encryption*, IDS [@] , VPNs [@]	1*, 2*, 3 [@] , 4 [@]	Traffic Analysis* Traffic Flow Analysis* Perimeter Protection [@] Authenticity [@] , Confidentiality [@]
Boundary Layer 1 (NMCI Connection to NIPRnet and SIPRnet)	Firewall, IDS, VPNs [#] , Link Encryption*	3, 4, 7 [#] , 1*, 2*,	Perimeter Protection, Traffic Analysis* Traffic Flow Analysis* Authenticity [#] , Confidentiality [#]
Boundary Layer 2 (Legacy Navy Networks)	Firewall, IDS, VPNs [#] , Link Encryption*	3, 4, 7 [#] , 1*, 2*	Perimeter Protection, Traffic Analysis* Traffic Flow Analysis* Authenticity [#] , Confidentiality [#]
Boundary Layer 3 (Communities of Interest)	VPNs, Firewall**, IDS**,	7, 3**, 4**	Authenticity, Confidentiality, Perimeter Protection**,
Boundary Layer 4 (Hosts and Servers)	Host-IDS, Anti-virus, Configuration Management, Smart Card ^{\$} , Email Encryption ^{\$} , Web Server Authentication ^{\$}	3, 4, 7 ^{\$}	Perimeter Protection, Authenticity ^{\$} , Confidentiality ^{\$}

Table 14 Boundary System Summary

Notes

* used only on the classified side of the boundary

@ used only on the unclassified side of the boundary

VPNs used only for applications that are not compatible with NMCI firewalls

** not used across all of the VPNs in boundary

\$ used only with the PKI implementation

This table highlights one of the problems with the deployed defensive systems within NMCI. The application of software

and hardware is inconsistent across the boundaries both internally and externally. While this by itself is not outside the norm when other enterprise networks are observed, there are potential problems with what these applications are intended to achieve and what that are best suited for. To see the potential problems we can examine them boundary by boundary.

a. Transport Boundary

The transport boundary is intended to move both classified and unclassified data between NMCI BANs/LANs via the vBNS or DISA services. The transport boundary employs intrusion detection systems and link encryption for security of its two internal layers. Referring to the Buddenberg matrix we can see that the intrusion systems are a mechanism to provide enclave security, while the link encryption is a method of preventing traffic analysis or traffic flow analysis. The link encryption is deployed on the classified side of the transport boundary only. The inverse is true for the intrusion detection systems which reside only on the unclassified portion of the transport boundary [RAY01]. VPNs are employed on the unclassified side of the transport boundary to connect the communities of interest that are separated geographically.

The communities of interest that reside within NMCI and utilize the unclassified portion of the transport boundary are likely the best protected of any within the network based on this arrangement. The combination of the VPN, transport layer IDS, and the host level IDS and configuration monitoring address the needs of confidentiality, authenticity, perimeter protection, and a secure pipe. There are gaps in the coverage, however. The

VPN is not an end to end encrypted path, beginning and ending at the VPN gateway within each NMCI LAN. The weakness to this arrangement is that confidentiality and authenticity of the data between the gateway and the destination host is not guaranteed. After reaching the gateway the data moves in the clear text from the gateway to the NMCI host. This gap is significant, given that the greatest threat is often from internal participants of the network [11]. This vulnerability could allow a member of the network the ability to exploit the data before it reaches the destination host. The assumption that supports this deployment of a VPN are that the NMCI BAN/LAN could not be compromised, the data is somehow otherwise unobtainable by an intrusion, or that any intrusion would be discovered. The historical experience of dealing with intrusions shows this last assumption to be utterly false.

Users who are not part of a COI do not benefit from the deployment of a VPN. For those users there is no application that provides them with any level of confidentiality, authenticity, non-repudiation or data integrity. The IDS at the network and host level serve the requirement for enclave security, but remainder of the transport boundary is essentially a commercial ISP service. While the network is isolated from other traffic, this is not a guarantee of security of the data while in transit. The data could be quite vulnerable to alteration in this environment. Viewed in an operational context, this practice accepts a significant risk.

Today, much of the logistical data sent within the defense message system is at the unclassified level. Data integrity is extremely important in this mission area and the potential disruption that could be caused by

changing relevant quantity, quality, or line item numbers on logistical messages is tremendous. Misdirected delivery of munitions, fuel or personnel to improper locations can result in units becoming combat ineffective. Corrections of these transactions while in process would require human intervention to perform and some element of chance or luck to detect before they have actually occurred. The prospect of catching these malicious acts before they produce a significant disruption in a high-volume, high-tempo environment is extremely low. The upshot of this is that even though the segment of the network is unclassified it requires an application or mechanism to support data integrity, authenticity, and non-repudiation to required by the mission. This need for authenticity is universally applicable to all traffic that rides within NMCI. There should be an authentication mechanism for all the official business transactions that occur.

On the classified layer of the transport boundary link encryption is employed to provide confidentiality and resistance to traffic analysis/traffic flow analysis while relying on the intrusion detection systems of the outer routers of the BAN and the host based IDS for perimeter protection of the network. This arrangement, when compared to the Buddenberg matrix constitutes a misalignment between the network requirement and applied solution. The encryption in the classified portion of the transport boundary is being used to address the need for authenticity, non-repudiation, and data integrity. Given this portion of the transport boundary moves classified data, it would seem logical to emphasize these characteristics for the data while in transit. This deployment doesn't allow for the authentication of the

originator of the data, the integrity of the data while in transit, or provide for the accountability of the receiver of the data after delivery. Accountability for who has accessed what data is of particular emphasis when handling classified printed media and there should be some mechanism for doing so within NMCI given the migration to non tangible medium for all types of information, both classified and otherwise.

b. Boundary Layer 1

Boundary Layer One is intended to provide connectivity between NMCI and the NIPRnet and SIPRnet. Like the other boundary layers, boundary layer one has both a classified and unclassified portion. Firewalls, IDS, and VPNs are used within the unclassified portion of boundary layer one to provide security. The network firewalls and IDS work in concert with the host level IDS to produce the enclave protection when an NMCI host utilizes this boundary for access to the NIPRnet. The VPNs employed within the unclassified side of the network are applied in a manner different than in the transport boundary discussed previously. The VPNs on the unclassified boundary layer one are there to support access to legacy applications that reside on the opposite side of the NMCI firewall from the user. These legacy applications are not or can not be made compliant with the current NMCI firewall policy. Rather than weaken the firewall on the NMCI BAN/LAN the choice was made to use the VPN to transport the legacy application through the firewall to give the users the required access. The weakness to this arrangement lies with the legacy application and the relative misuse of the VPN in this role.

The legacy application is likely non-conforming with firewall policy because of identified exploitable insecurities. VPNs are intended to provide a degree of authenticity and confidentiality, not necessarily to address an enclave security issue as these legacy applications appear to represent. The VPN use in this case provides a point of entry for the insecurity that could not be obtained otherwise. If the legacy application has been compromised, the VPN is simply providing the intruder a door to bypass some of the enclave protection the boundary is intended to provide. There is a cofferdam configuration that is built into VPN employment that allows the datagrams to be decrypted, examined for defect by IDS, re-floated via another VPN device, and then forwarded to the destination gateway.

This arrangement could catch identified vulnerabilities, but again the vulnerability must be known to exist for it to be effective. This employment of the VPN in this case also suffers from the downstream lack of encryption to the destination. The VPN carries the data in encrypted form only as far as the VPN gateway in the BAN/LAN and from there it moves in the clear. As in the use in within the transport boundary there is the chance of compromise from within the NMCI network while the data is moving between the VPN gateway and the NMCI host.

The classified layer of boundary one employs both link encryption and VPNs for security and provides connectivity between NMCI and the SIPRnet. It is assumed that there exist within this boundary some communities of interest (COIs), and so the employment of the VPNs in concert with the link encryption would provide those members with an effective level of authenticity and

confidentiality, in addition to the benefits of the enclave security. Outside of a COI, however, there is no identifiable mechanism for providing assurance of authenticity or non-repudiation within this layer of the network. The sophistication of the Type One encryption presents a significant barrier to entry; however it is not principally intended to provide authenticity to the data. The assumption appears to be that this sophisticated encryption effectively does just that. Encryption has been shown to have its limits and reliance completely upon it for security of classified information accepts less security than is placed on other forms of classified data. In any event, the use of encryption to provide for authenticity is a misalignment of the requirement and the solution within the Buddenberg Matrix.

It should be noted that greater demands for authentication and data integrity are placed upon the examination of hard copy classified media. For some classification levels and some types of classified material, viewing and handling it requires two persons be present at all times. Two-person integrity is cumbersome and may be equally so on a network, but allowing the information to move so freely with a networked environment without some aspect of traceability and data integrity check is an insecurity we do not accept with hard copy classified material. Given the ease of movement of digital data there should be some application for achieving authenticity to place network access on the same footing as hard copy access. Simply encrypting the data for transit while providing the end users with the decrypted text leaves open the issues of who viewed the material and at

what time and place, and did all of the data return without being duplicated?

c. Boundary Layer 2

The construction of Boundary Layer Two is essentially identical to that of boundary layer one. Firewalls, IDS, and VPNs are deployed in the same manner as in boundary layer one. Only the unclassified portion of boundary layer two is described in any detail. The classified portion is only differentiated in the documentation by the mention of the use of type one encryption "where needed" [RAY01]. There are a multiple of access configurations that could be present in the unclassified portion of boundary layer two and the final definition of those has yet to be determined. As part of the NMCI SSA five different representative scenarios are presented as possible solutions to this problem. The five scenarios are described as follows [RAY01];

Scenario 1: NMCI Hosted DON Legacy Server

Scenario 2: Non-NMCI DON Legacy Server

Scenario 3: Joint Non-DON Legacy Server

Scenario 4: Joint Non-DON Hosted Server, Replicated

Scenario 5: Joint Non-DON Hosted Server with Non-DON VPN

Viewed in terms of the use of VPNs, scenarios one and three employ the VPN as a mechanism for access to a legacy application that is resident on the legacy network and is not firewall compliant. In each of these scenarios the VPN is being employment as an access tool for the NMCI user to reach the legacy application. The difficulty here is the same as was seen in boundary layer one. The VPN is being

employed as an enclave defensive system when its purpose is to provide confidentiality and authenticity, not a protection against exploits that are buried within the legacy application itself. Whatever burdens the legacy application possesses are brought into the NMCI environment at the BAN/LAN level.

Scenario Five employs the use of a network to network VPN, combined with type 1 encryption to transfer data via the SIPRnet, between another service's classified BAN/LAN and an NMCI BAN/LAN. Upon entering the NMCI BAN/LAN the joint VPN is routed through a specific joint service VPN gateway and then into the NMCI environment to the end user. The VPN provides service only when the data is on the WAN. This effectively hides the information while riding on the SIPRnet (which is also employs type one encryption) between the two classified BANs. This could be used to prevent the disclosure of sensitive data to SIPRnet members that do not possess the "need to know", but again the data moves unencrypted within the NMCI classified BAN to the end user. If the data is so sensitive as to require limited exposure within a classified network, then the data would likely require equivalent confidentiality end to end. The NMCI classified BAN however, doesn't provide this level of confidentiality and so the need for confidentiality and authenticity is not addressed to its fullest in a situation that may require it.

Scenario Four is a replica of the scenario five deployed for access through the unclassified NIPRnet. This boundary uses a network to network VPN to move the data from the legacy server to the NMCI BAN via boundary layer one, and then a single sided VPN moves it to the NMCI BAN and the end user. The legacy application resides in the

DMZ external to NMCI and behind a firewall operated by another service. The employment of the VPN in this manner implies the application may not be firewall compliant though this is not specifically stated. This implication is supported by the use of a single sided VPN to move the data past the NMCI firewall. The VPN terminates at a gateway and then flows to the end user over the NMCI BAN. Given the data is unclassified and the VPN is single sided the primary purpose must be to avert the conflict between the firewall and the application providing the user access.

Scenario Two is the most basic of the five and uses the Boundary Two firewall and IDS as the mechanisms for defending the network. This is the conventional enclave arrangement and as such does little to provide the user with any level of confidentiality or authenticity of the data. The firewalls and IDS are the static defenses that are relied upon for protection. What is unclear about this boundary protection is that the legacy applications accessed are identified as not firewall compliant, so it is uncertain how these identified non-compliant applications will be accessed through that same firewall.

d. Boundary Layer 3

Boundary Layer Three is designed to give separation between specific communities of interest that lie with NMCI. There are four communities of interest identified within NMCI and they are mapped on the table below.

	Group	Virtual LAN (VLAN)	Shared VPN Gateway	Dedicated VPN Gateway	IDS	Firewall
Sensitive (A)	X	X ¹				
Highly Sensitive Distributed (B)	X	X	X	X ²		
Highly Sensitive Co-Located (C)	X	X			X	X ³
Isolated (D)				X	X	

Table 15 NMCI COIs

Notes: 1. To limit network access to a private server
2. To protect private server or enclave with its own LAN
3. If required

Communities of Interest A through C employ a virtual LAN in concert with the host based system for enclave security. Because of the dispersed nature of COI B a VPN is deployed to provide confidentiality and authenticity. COIs C and D benefit from the use of an IDS and either a dedicated firewall or VPN as required. This configuration is likely the most effective within NMCI, but is limited in its application because of the comparative small size of the groups involved. The VLAN application and host defenses provide reasonable enclave security, while the use of VPNs serves the needs of confidentiality and authenticity to the users. The weakness existing in the VPN encryption is not truly end to end, providing an opening to an internal intruder. The relative small size of the COIs may mitigate this to a degree, however it is not as effective as true end to end encryption between users.

e. Boundary Layer 4

Boundary Layer Four is the host server level of the NMCI defense. This layer utilizes host based IDS, anti-virus applications, and configuration management. In conjunction with these applications there will be a Public Key Infrastructure (PKI) implementation. Unfortunately this will not be implemented as part of the roll out and so the true date for PKI operation is presently undetermined. Without the PKI implementation the boundary four protections provide enclave protection only. Authenticity and confidentiality are provided via the password protections which is a conflict between the requirements and solution with the Buddenberg matrix. The VPN clients that are deployed provide this confidentiality when used, but we have seen that this application only goes as far as the gateway and provides only limited capability. The IDS and antivirus applications are reactive in nature and only as effective as the management tools that keep them up to date.

While effective at protecting the box, these are of little value to the network in total and do not address the full spectrum of security requirements. What is needed for NMCI is a method of security that effectively provides security for the data that is being moved in addition to the security of the boxes and pipes that constitute the network.

f. Security Summary

While the security applications deployed within NMCI are effective at providing a degree of enclave security, there are significant gaps in the coverage they provide. There are obvious misalignments in the system

requirements and the applied solutions throughout the boundary layer system.

The transport boundary relies heavily upon the use of link encryption to provide authenticity and confidentiality of the data as it moves on the classified portion of the boundary. The VPNs deployed on the transport boundary are gateway to gateway and do not protect the data all the way to the end user. In Boundary Layer One the VPNs are employed as a mechanism for accessing legacy applications that are not compliant with the NMCI firewall policy. Boundary Layer One also employs link encryption as a means of providing authenticity and confidentiality. In Boundary Layer Two, three of the five offered scenarios employ VPNs in the same manner as boundary layer one, to provide access to legacy applications. Boundary Layer Three is one of the better arrangements within NMCI, but here again is the limitation that the VPNs employed cover the data only as far as the gateway within the BAN or LAN. At Boundary Layer Four the defensive mechanisms support only the enclave defense and do not provide any level of protection for the data contained therein. Each of these solutions constitutes a misalignment of the applied solution and the network requirement when mapped onto the Buddenberg matrix.

In addition to these misalignments, there appears to be a lack of consideration of the threat posed by an internal network member. Even in the portions of the respective boundaries where the solutions and objects are reasonably aligned, there are gaps that could be exploited by an internal intruder. The VPN devices provide protection only as far as the gateway and not all the way to the end user. These gaps offer an opportunity to

members of the network to sniff or capture the traffic as it flows internal to the BAN or LAN. An intruder would have to defeat some of the internal network defenses to deploy his own exploit, but this has been done successfully in the past and is likely to occur again in the future. Effectively securing the data could negate this effort or at least make its execution of much less value.

There may be an assumption that internal compromise is unlikely, but when considering the historical cases of compromises to Navy security in other areas, the internal members have proven to be the most damaging. The John Walker family espionage case proved that significant and long term intrusions can go unnoticed or unchecked and produce tremendous damage. Without more effectively addressing this weakness NMCI could suffer a similar compromise. This is of particular concern given that NMCI is obligated to support the numerous legacy applications and networks that are in the Navy inventory today.

Legacy applications represent a distinct challenge during the initial transition into the NMCI environment because many of them were created before network security was given the significance it possesses now. If there is a weakest link in the network chain it is these legacy networks and applications that are based upon operating systems software that possess demonstrated compromises. Consequently, inclusion of these legacy applications and networks poses a significant risk to NMCI in its initial phase of operation. This is not meant to suggest that this issue has been ignored; rather it is obvious when looking at the NMCI security boundary system that this issue has been addressed.

Stepping back from the specifics of the individual boundary definitions their larger purpose can be seen. Of the five defined boundaries within NMCI, Boundary Layers One and Two exist primarily to provide access to legacy applications or networks that reside within the Navy and Marine Corps. Contemporary networks are difficult enough to maintain security within, and these legacy systems makes the job for NMCI that much more difficult. Some of the difficulty lies in the methodology that is used in the application of that security. Securing the network pipe and the network box, as is attempted within NMCI, has limited effectiveness without greater attempts at securing the data that exists within it.

Securing the data within NMCI is dependent upon the inclusion of an object level security approach to the problem of securing NMCI as a whole. Looking back at the OSI layer model and the Buddenberg Matrix, and then comparing the layers of the NMCI enclave we can see that nearly all of the NMCI protective features are placed at OSI layers 1, 2, 3 and 4.

Service	OSI Layer
Confidentiality	1, 2, 3, 4, 7
Authentication	3, 4, 7
Integrity	3, 4, 7
Access Control	3, 4, 7
Non-Repudiation	7

Table 16 ISO 7498-2 Layers

VPNs address the layer seven requirements incompletely because of their configuration within the network. In Mahanian terms this would be like escorting

your merchant ships only seventy-five percent of the way to port and hoping the enemy raiders just don't show up in that other twenty-five percent. The data require a convoy system for protection while in transit. The data needs protection port to port and node to node for the network to be successful in its mission. Object level security is the only method of addressing all of the OSI layer security requirements at the same time and at the same OSI layer. Application of object level security also carries with it additional benefits to the network as a whole.

Employing object-level security requirements will alleviate the need for the use and management of some of the encryption and VPNs employed within NMCI now. With object level security the network pipe can be come more generic because the security functions are being provided at the application level and not the transport level. This eliminates the need for widespread encryption throughout the network and allows it to be employed only where there is a need to protect the network pipe. This gives the network greater flexibility, reduces the management burden and cost, while providing effective security through the alignment of the requirement and the applied solution. The need to secure the network box and transport layer may not be completely eliminated, however network security is greatly enhanced through the application of object level security by effectively addressing all of the network security needs on a common OSI layer, making its application and effectiveness much less problematic.

H. QUALITY OF SERVICE

An effective quality of service implementation is essential to the network being able to perform its mission tasks during periods of limited availability or restricted bandwidth as the result of compromise or damage. Differential services within the network routing structure are what give a network the ability to cope with these problems in a logical and predetermined manner than supports the networks primary mission functions. The functionality provided by differential services is applicable down to the BAN and LAN level, but is most relevant at the WAN level and above that connect the various NMCI enclaves deployed in the U.S. and around the globe. This translates to the vBNS and DISA services that constitute the transport boundary within NMCI. Differentiation at this level is most relevant because the communicating parties are geographically isolated from one another by large distances. In addition, the vBNS and DISA will likely be handling the largest volume of prioritized traffic when compared to any individual WAN or metropolitan area network within NMCI. This requirement is even more critical when considering the need to provide a global connection to the Fleet.

NMCI is tasked to provide connectivity in cooperation with a global information grid. This implies the need to connect to the Fleet underway, likely through a radio wide area network (radio-WAN). The radio-wan will likely be a bottleneck for communications with the Fleet. Consequently the radio-WAN will need an effective means of applying differential services to its traffic. NMCI, operating in cooperation with that radio-WAN, will need to be able to

respect the differential service scheme employed within the radio-WAN as it moves the incoming Fleet traffic.

The quality of service implemented within NMCI is a designed to provide specific applications with reserved bandwidth, controlled jitter, latency and packet loss. The vBNS system operated by MCI Worldcom utilizes multi-protocol label switching (MPLS) to perform traffic engineering within the backbone of the transport boundary of NMCI. The MPLS used by Worldcom employs resource reservation protocol (RSVP) to develop the quality of service guarantees for various application flows. The key feature of the MPLS is its ability to provide label switched paths (LSP) which are similar to permanent virtual circuits (PVC).

The MPLS works by measuring the available resources (bandwidth) and then allocating them to LSP tunnels, which are explicit flow paths from ingress to egress of the vBNS structure. MPLS traffic engineering routes traffic flows across a network based on the resources the traffic flow requires and the resources available in the network. An interior gateway protocol measures the flow within the tunnels and the demand for service and then dynamically reconfigures the tunnels to fit the required load. If any particular stream exceeds the capacity of a LSP tunnel, multiple tunnels will be allocated to the same ingress and egress points to carry the traffic. This application allows the vBNS backbone to support a high use of transmission capacity while being very resilient, so that it can withstand link or node failures. This gives NMCI a higher degree of availability of the backbone than might otherwise be available, but does not fully address the need for differential service.

The MPLS system allows the network to achieve greater connectivity during restricted bandwidth because it dynamically matches the available bandwidth to the LSP tunnels. The problem with this application is that it meets the needs of connectivity based on the available bandwidth without restricting access to that bandwidth on any other criteria. The LSP tunnels are determined by the destination of the packet and not by the mission function or originator of the packet being sent. Consequently, bandwidth is likely allocated to functions that may or may not be considered mission essential during times of restricted bandwidth. There needs to be an effective mechanism for discriminating between high bandwidth demand applications based on mission function. The latest AFRTS (Armed Forces Radio and Television System) release and real time video from a remotely operated vehicle would demand similar bandwidth. But how does the network determine who gets the available bandwidth, Arnold Schwarzenegger or Osama Bin Laden? MPLS, while enhancing the availability of the network, does not address the natural tension that exists between the need of all to be connected and the need to maintain mission essential functions. As bandwidth is reduced, the internal gateway protocol will attempt to keep as many tunnels open to as many locations as possible, allocating bandwidth reductions in line with the existing ratios. While this is desirable, at some point there must be a decision to allow some members availability to fall out in favor of other more relevant and mission essential organizations or units.

This QoS application discriminates based on destination IP address and not the mission functions contained within the network or possessed by specific

application flows. What is needed is a protocol that will allocate the available bandwidth in a dynamic manner based upon the mission precedence of a given demand for traffic flow. The highest mission precedence should be allocated the required available bandwidth first. Lower precedence traffic should be required to wait in a queue until it possesses adequate precedence compared to competing traffic and can be forwarded.

Implicit with this idea of precedence is that at some point available bandwidth to some users will effectively become zero based on the assigned or designated precedence of their transmissions. This can be likened to the imposition of "minimize" that exists within the defense message system. The imposition of "minimize" requires users desiring to transmit a message to a specified destination that is under "minimize" to obtain a precedence of a certain level. This is intended to cull out much of the non-mission essential traffic that would otherwise be sent. Differential services follow this same pattern, only in reverse.

Under restricted bandwidth conditions users would be restricted or prohibited from transmitting on the network based upon the precedence assigned to their traffic. This imposition would likely be based upon user identification within a particular command, as is the message release authority for the defense messaging system. If adequate granularity could be achieved the optimal solution to this problem would be to base it on the specific application being used. Placing the deterministic factors for differential service at the application layer would permit all users access to the network if the precedence of the traffic deemed it appropriate. This keeps some degree of

connectivity to each organization at each level. Placing the deterministic characteristics at the application level also aligns the need for the performance of the differential service with the security requirements as expressed within the Buddenberg matrix.

Another potential limitation of this QoS application is that it relies on RSVP to create tunnels within the network. The potential problem is the unintentional development of logical single points of failure within the network through the reservation of resources. This could be the result of the attempt to remain connected to as many nodes as possible without regard to effective bandwidth or availability.

While the implementation of quality of service within the NMCI vBNS system is superior to the best effort service, it falls short of providing adequate ability to differentiate between mission essential and non-mission essential service application flows. The MPLS will enhance the availability of the network, which improves its survivability, but it doesn't adequately enhance the performance of the mission essential functions within NMCI to the point of making it a survivable network.

I. NMCI SECURITY ARCHITECTURE CONCLUSION

The NMCI security architecture is likely superior to many enterprise implementations that are operating today. It possesses significant improvements over the basic concept of network security architecture through its employment of multiple internal layers. NMCI does not, however, possess the requisite characteristics necessary to make it a survivable system. There are lapses or gaps that need to be addressed within the structure.

The enclave security strategy is arguably an incomplete practice, particularly in terms of the needs of a modern network intended to perform a military mission. Enterprise implementations in the private sector are comparable, but the fact that is overlooked by NMCI is that even these business organizations have a central mission focus to their network. NMCI is attempting to perform all missions for all masters and this is reflected in the enclave strategy. Efforts were made to protect everything. Unfortunately there were no efforts made to enhance any of the specific mission capabilities of the network. Alfred Thayer Mahan's theories suggest we need to be offensively minded if we desire to dominate the battle space. The network environment is the definitive battle space of the future, if not the present. With NMCI we have chosen to be purely defensive, attempting to deny the denier. The network security strategy within NMCI needs to change to reflect the needs of network centric warfare, not network centric defense.

The focus on network-centric defense can be seen in the security systems within NMCI. The boundary layers that constitute the network reflect the dependence upon contemporary network defensive strategy. The tools are essentially reactive and possess very limited ability to adapt on their own. The defensive mechanisms are largely constructed upon historical data and don't adequately address the implications of future attackers.

The availability provisioned by the relevant SLAs do not fully address the standards of dual threaded capability. While not a certainty, the contractual agreements do not require the service provider to specifically meet these levels in the design or

construction of their systems. The availability levels established are below the dual redundancy demanded in other major systems employed in the warfighting function of the Navy and Marine Corps.

The quality of service implementation within NMCI is certainly better than could be had in most other networks. The effect is to enhance the availability levels of the NMCI backbone, which is a benefit based upon the examination of the contractual levels obtained. Unfortunately, as good as this application is, it falls short of the level of differential services required to protect the networks mission essential functions.

At the end of the day the NMCI architecture, while measurably more effective than previous designs, fails to meet the requirements of survivability because of its inability to protect and preserve the networks mission essential functions. Unfortunately for the Navy and Marine Corps this failing comes at great cost, and growing risk.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

A. INTRODUCTION

The analysis of the Navy Marine Corps Intranet undertaken in this thesis evaluated, at a relatively high level of abstraction, the mission, mission functions, and architecture of the network to determine if the design was consistent with the concepts of the network survivability analysis method. In the process of this evaluation, it became evident that concepts of survivability can be better defined by tying them to the contemporary framework of networks. Doing so allows for an easier translation for those who are less familiar with the original documents and places them into a context more suitable to the applications and hardware that will perform the necessary network tasks.

The original work on network survivability identified the key properties as recognition, resistance, recovery, and adaptation. When these definitions are mapped onto the more conventional terminology used in the discussion of networks it can be seen that there is some overlap among them. The characteristics of a survivable network can be better parsed so as to eliminate the overlap and more clearly separate the qualities and their relevant required functions.

I believe the requisite survivability attributes that a network should possess are better defined as mission, availability, security, and quality of service. These characteristics map more explicitly to the software and hardware applications that comprise a network and the functions they perform both individually and collectively.

It is also important to understand that these four characteristics and the overall survivability of a network are interdependent. A network needs to possess all four of these characteristics to be survivable and the failure of any one places the overall network survivability in doubt. Of the four, however, the characteristic of mission definition is the most significant.

The network mission definition is the most important characteristic to network survivability because it directly influences the implementation of hardware and software that combine to create the other three characteristics within the network. The opposite can not be said to occur. The application of hardware and software within a network can not combine to produce a mission definition for a network because the mission definition is derived from the network designer's/user's intent. An examination of the resident hardware and software and the resultant internal characteristics of the network can infer what mission(s) is (are) relevant to the network or what mission functions it could perform, but it can not derive intent. Therefore, networks that possess a poorly defined, nebulous or too broadly characterized mission definition possess a fundamental flaw that likely inhibits their ability to achieve survivability. Such is the case with the Navy and Marine Corps Intranet.

B. NMCI MISSION DEFINITION

Navy and Marine Corps Intranet survivability is fundamentally flawed because of the lack of a clearly defined mission function. The network is required to support the core functions for business, scientific, research, computational activities, and warfighting.

Essentially NMCI is required to perform every mission for every man. Consequently, determining the mission essential functions of the network was impossible, since all missions and all functions were given equivalent footing within the network architecture. The examination then became one to determine what mission the network was best capable of performing based upon the established architecture.

In broad terms, I defined the most demanding mission for NMCI as force projection, disaggregated this mission into the two mission essential functional flows of logistics and readiness, in that order of precedence. NMCI can begin its progression toward survivability by adopting the force projection mission definition and developing the mission essential functional flows of logistics and readiness I have identified in this thesis. By doing so, NMCI's mission definition will become aligned with a core warfighting mission requirement of the Navy and Marine Corps.

C. NMCI LEGACY AND TRANSITION

Given the number and varied composition of software operated by the DoN and the essentially finite amount of funding available there will likely always be something that falls under the legacy definition. The existence of legacy systems and the requirement to transition them will then be an ongoing issue for NMCI for some time to come. Legacy is a double-edged sword for NMCI. These systems hold vital, valuable data to the enterprise that needs to be preserved and passed forward so that NMCI can effectively perform its assigned mission. Legacy systems are also a threat to NMCI because they offer a weak point for exploitation, potentially becoming legitimized trojan

horses used to violate an otherwise more secure environment. The question for NMCI then is not if it must accommodate legacy, but rather how it must accommodate legacy in a manner that is consistent with NMCI's mission definition.

The selection of legacy systems for transition should be based upon the contribution the individual application makes to the primary mission function of NMCI. Based upon the force projection mission definition I offered for NMCI, the legacy applications chosen first for transition into the NMCI environment (from the list of approximately 37,000) should be those that support the essential logistical functions of the Navy and Marine Corps. The next applications selected should be those that support the primary readiness and training functions of the Navy and Marine Corps. Transition of applications in this manner balances the risk and benefits to the network in a logical, mission based manner. Since all legacy applications represent a potential threat, it is logical then to bring only those that enhance the network's ability to support the Navy and Marine Corps combat capability. Transitioning any application that does not meet this requirement constitutes the acceptance of unnecessarily greater risk to the network.

The organization of the transition of the Navy and Marine enterprise should be guided in the same manner. The transition into the NMCI environment has been guided by the traditional organizational lines that dominate large organizational thinking. The failing of this methodology is that the transition of any single unit or staff is not effectively complete until the entire enterprise has transitioned. This is because the benefits of increased

functionality are best realized only on the enterprise level. The transition into NMCI should be functionally based upon NMCI's mission definition and the transition executed on an enterprise level. By transitioning in this manner the functionality provided by the individual application is achieved force-wide in a single event and the requirements for remedial work are reduced significantly. To guide the execution of the transition and the implementation of the legacy software I have offered a practice from the software design community known as spiral development.

The application of the spiral design methodology to the NMCI transition will permit functionality to be implemented force wide while allowing an incremental approach to the execution of the entire process. The advantage of this approach is two-fold. First, employing the spiral method to the transition into the NMCI environment will produce functionality across the entire force in each iteration. If functions are transitioned instead of command organizations, then all the members that employ that functionality gain.

Assuming the functions are transitioned in priority based upon mission function, the entire force gains because the functionality is connected to the Navy and Marine Corps core mission requirements. Second, the spiral method allows for the development of the mission essential functions for the network, a critical requirement for the achievement of survivability. The definition of NMCI's mission essential services can then be used to guide the application of availability, security, and quality of service implementations within the network architecture in the effort to achieve network survivability.

D. NMCI SECURITY ARCHITECTURE

The NMCI security architecture was examined for the characteristics of availability, security, and quality of service, the elemental requirements of network survivability. How effectively these elements are implemented within the network determines if NMCI is a survivable network.

The levels of availability provisioned within NMCI are reflective of other enterprise network implementations. Unfortunately, this is not entirely adequate to meet the mission requirements for a survivable network or for a network that entails a military mission function. The effective availability for the network is less than the dual redundant capability that the application of survivability reflects in other types of complex systems. Ships and aircraft possess this redundancy at some level, the focus of which is to keep the ship afloat or the aircraft in flying. The same can not be said for NMCI. The levels of availability within NMCI do not reflect the mission essential status the network has been given.

Quality of service is implemented within the network, but not in manner that supports survivability requirements. Application flows are given priority based upon their function without consideration to the network mission function of the flow or the precedence of the originator. For a network to be survivable it must provide those mission essential functions while under duress. To meet this requirement NMCI needs a prioritization of internal functions and originators and a mechanism for traffic differentiation. The quality of service within NMCI doesn't present this capability and as such does not fully meet the needs for a survivable network.

The security mechanisms within NMCI are as current as any network system in operation. Unfortunately, they are rooted in a defensive network security strategy that may be fundamentally flawed. At a fundamental level, networks are designed and built to move information while network defensive strategy hopes to protect the network by somehow restricting the movement. What is needed is an offensive network strategy that emphasizes the ability of the network to move the data to the location it desires in spite of the success of any attack against it. Network survivability is the strategy that moves NMCI in that direction.

The applications that perform the security function are likely effective in many areas, but there exist misapplications of technology and gaps in coverage that present opportunities for exploitation. More importantly, the security applications within NMCI are potentially sub-optimal because they seek to secure the network boxes and pipes, ignoring the importance of the information that travels within them. Shifting the focus to securing the data within NMCI could produce a more cohesive while less coupled network, significantly enhancing the survivability of NMCI.

To achieve this, the authentication and confidentiality mechanisms must be internalized to the specific applications. Embedding these functions into the individual applications secures the data being transmitted and allows the application to be completely indifferent to the physical network that the data is traveling on.

E. NMCI SURVIVABILITY CONCLUSION

To summarize the work conducted in this thesis, NMCI is not a survivable network for the following four reasons;

1. Lack of a clearly defined mission or missions
2. Availability (A_0) that is less than needed to ensure the retention of full mission capability.
3. The quality of service implemented does not provide for application of differential service of network traffic.
4. The security mechanisms employed do not ensure the security of the data within NMCI.

While all of these failings are significant, three of the four are essentially technical problems. Technical problems have not proven to be insurmountable. The most difficult problem faced by NMCI is the lack of a clear mission definition, and until this is determined, no amount of technical solutions will produce a survivable structure for NMCI.

F. RECOMMENDATIONS

If the U.S. Armed Forces are to transition to network centric warfare, then one among them must make the first leap into the waters of cyberspace. The Navy and Marine Corps team have chosen to be the first with the advent of the Navy and Marine Corps Intranet. The problem being faced now is how to make this system viable for use in a military application. There are several things that should be done to correct the drift of NMCI and lay it on a course

that will produce direct, tangible results that support the Services and their warfighting mission.

First, NMCI should assume as its primary mission force projection, disaggregated into the logistical and readiness mission essential function flows. This will begin the process of segregation of functions within the network and among the legacy applications awaiting transition that is necessary for the construction of an adequate quality of service application for the network.

Second, NMCI availability should be contracted to the level of four 9's (.9999) of availability throughout. This measure should be taken from network host to network host. Given the series-dependent nature of NMCI, this level of service makes the requirements for SLAs addressing availability at other levels much less significant if not irrelevant. If the availability measured host-to-host is at four 9's then the cooperating network segments then must be higher than that. This is likely the easiest fix to be performed. The addition of hardware and planning of additional alternate routes is not a significantly complex task in most cases.

Third, NMCI must implement a quality of service control that will provide granularity at least equivalent to the existing JANAP 128 notion. While crude in comparison to the potential capability of a QoS application, it will require NMCI to meet the standard applied within our existing radio networks.

Fourth, NMCI must implement an object level security strategy for the entire network. Object level security will enhance overall security of the data while decoupling

the security mechanisms from the physical network. Decoupling security from the physical network provides NMCI with the greatest flexibility operationally and in terms of the overall NMCI business plan fee for service.

Lastly, NMCI needs to sponsor good housekeeping rules to guide the transition of legacy applications both now and in the future. Foremost among these rules should be the requirement for object level security and the need for alignment of the individual applications function with the overall mission of the network.

In total these recommendations constitute threshold requirements for networks performing in military applications. These requirements can be expressed in terms of availability, security, and quality of service and are mission dependent. Using the broad mission definitions from chapter two (Administration, Force Projection, Battle Management) we can see how the requirements increase with the desired mission function.

	Administration	Force Projection	Battle Management
Availability	.999	.9999	>.9999
Security	Enclave	Enclave + Object Level Security (sender to receiver)	> Enclave + Object Level ?
Quality of Service	Best effort	JANAP 128 standards as a minimum	> JANAP 128

Table 17 Mission Area Threshold Requirements

It is beyond the scope of this thesis to refine these characteristics further. It has established what I believe to be the threshold requirements for the utilization of networks in a military mission. A network requirement that falls to the right of the force projection mission will likely require a greater application of these qualities to meet the mission need. This is due in large part to the greater demands of the specific mission area. A requirement that falls to the left of force projection mission area begs the question of why construct a Service specific intranet in the first place.

Adopting these recommendations will place NMCI on the road to survivability without undue effort compared to what has taken place so far. Adopting the mission function and reorganizing the legacy transition are a matter of emphasis. The increase in availability levels are the simple application of additional hardware where needed. The quality of service implementation is the most difficult, but is likely solvable in a short period given the comparatively crude sieve desired for network traffic. Force projection is likely not the proverbial "killer app" that the Navy and Marine Corps sought with the inception of NMCI. It is in my opinion a near certainty that, if the rather mundane but vital functions of logistics and readiness cannot be mastered in the networked environment, then we have little hope of ever obtaining any viable capability close to the concept of network-centric warfare.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A

This appendix contains the lecture notes of Professor Rex Buddenberg and addresses the subject of availability of a networked system. These notes are the work of Professor Buddenberg and he is responsible for their content. This appendix is included for the readers benefit and to gain a better understanding of how availability applies to networked systems.

Availability (and survivability)

"Amateurs talk about tactics; professionals discuss logistics"

.....sign on John Lehman's desk[1]

Network analog: Amateurs talk about bits per second; professionals discuss availability.

Buddenberg/Apr95

I. Why is the subject of availability important in networking?

II. How do we do the arithmetic?

III. Putting it into some perspective.

I. Why is it important to know something about availability?

In a network centric set of information systems, sooner or later, enough of the operation uses the network that it can be viewed as mission critical[2]. Indeed, a network centric approach to systems engineering tends to result in several information systems (which may or may not be interrelated) all using the same network

substrate. A failure of that networking substrate brings many parts of the overall operation to a halt -- an increasingly unacceptable circumstance. The good side of this state of multiple systems using the same network is that it becomes increasingly economical to make the investments in high availability in the network because you only have to do them once rather than once for each information system. The other piece of good news is that all the tools required to build high availability networks are available off the commercial market. Indeed, by distributing information systems across a highly available network, we can build more survivable information systems through distribution than by using more traditional combat system engineering methods. (Unfortunately this is a subject that I've never found treated in networking textbooks; the quantitative parts here are mined from a reliability engineering text).

II. How do we do the arithmetic?

Ao is the engineering symbol for operational availability. It is usually expressed as a percentage and is defined as: up time / total time[3]. For example, the telephone company may quote you a 99.7% Ao, assuming you can find someone in the phone company who will tell you. Given 8640 seconds in a day, this means that for 8614 seconds the phone system will be responsive to you and for 26 seconds it won't -- on the average.

$$0.998 = \frac{\text{up time}}{8640}$$

Note that (up time) can be represented as (total time - down time) so we can solve for down time directly if we wish:

$$0.998 = \frac{8640 - \text{down time}}{8640}$$

Usually it makes much more sense from a requirements point of view to specify tolerable down time.

In complex networks we have several components strung together:

Serial arrangements.

We can either obtain availability figures for components from the suppliers or we can estimate them from experience. In this case let's assume:

Ao of WAN as 99.7%

Ao of router as 99.9%

Ao of the LAN and end systems (collectively) as 99%

Since the three components are wired in series, the Ao for the system as a whole is the product of the three component values:

$$Ao = 0.997 * 0.999 * 0.99 = 0.986$$

Figured over a month, this Ao figures as 605 minutes of down time:

$$0.986 = \frac{43200 - \text{down time}}{43200}$$

Can your network stand 5 hours of down time per month? (To be fair, note that the downtime is as likely to occur at night as during working hours). While it has been common practice in many DoD systems to attempt to improve readiness rates by increasing the Ao of the components, one can see that there are severe limits. And the limits are prohibitive in solid state systems such as current technology networks -- we have to replicate components and solve our Ao problems through redundancy.

Three principles of high availability engineering:

- eliminate single points of failure (often called common-cause failures)
- provide reliable crossover (from primary to backup)
- promptly detect failures upon occurrence

The rest of this section addresses the first of these principles. The second is neatly and inherently handled by the TCP/IP protocol stack for internetworks and by FDDI ring-wrap in LANs. The third is a core function of network management systems. These threads are taken up on those lessons.

Multiple-threaded systems.

Let's redraw our network layout to include two network installations:

Many networks grow like topsy so it's quite frequently that we find a pair (or more) of independent network installations [4]. In the stovepipe configuration above, both systems will exhibit the single-threaded Ao characteristics we've penciled out. But, let's twiddle a bit:

- let's assume that the larger WAN has multiple altroutes within it.

- bring the connectivity into the command/building/facility through two different central offices and through two different cable trenches. (Remember this

need when we talk about radio-WANs)

- cross-connect the routers (i.e. campus backbone)

- don't do something stupid like putting both routers on the same UPS or in the same wiring closet (different buildings makes sense)

- if the LANs are compatible, cross-connect them (one way is to add a bridge).

If we can reach a situation where one line failure can be compensated by the other, one router failure can be compensated by the other and component failures in the LAN can be compensated by redundant workstations and LAN cabling, then we've reached a point where recalculation of the arithmetic makes sense.

If the line Ao remains 99.7%, then the expected probability of failure is 0.3% or 0.003. Since we now have two lines, either of which being up represents success, then failure is represented by both lines being down: 0.003

* 0.003 or 0.000009. This means that the two lines working together have a combined Ao of .999991[5]. If we had a third line to contribute to the cause, and could still maintain independence of mode of failure, then we cook an Ao of 0.9999997!

Procedure:

- find probability of failure (1-Ao).
- multiply the probabilities of failure for all parallel systems.
- convert back to Ao by subtracting from 1 again.

Perform this procedure for each module: line, router, LAN. Given the hypothetical numbers we're using, and a simple duplication of the system, we get Ao for the pair of routers of six 9s and for the combined LAN assembly of four 9s.

- now multiply the three Ao values just as before:

$$Ao = 0.99999 * 0.999999 * 0.9999 = 0.999889$$

And you can then refigure predicted mean down time:

$$0.999889 = \frac{43200 - \text{down time}}{43200}$$

And get about 5 minutes of down time per month. This is a pretty dramatic shift -- at pretty nominal cost.

III. Perspective. What does all this arithmetic mean?

1. To the requirements setter.
2. To the network manager.

Requirements setter. First of all, make sure the requirements setter specifies an Ao requirement or objective. Of all the requirements documents I dealt with in 6 years in the business in CGHQ, I only had one that specified an Ao value -- and I'd ghostwritten it. In general, if you ask an operator what his availability requirement -- Ao = ____ (fill in the blank) -- is, he won't give you anything meaningful. Ask the question in terms of tolerable down time and then do the arithmetic.

Expect the requirements setters to lowball the availability requirement. This happens for several reasons:

- the sponsor is only worrying about his application on the network. If you look at the aggregate of several information systems residing on the network, the real availability requirement will be higher.

- a sponsor will often lowball because he's trying to chisel costs. Seems to be a natural human tendency to deny this requirement.

Network manager. Fortunately, with current internetworking technology, if we do a decent modularization job, these shortcomings can almost always be fixed later at fairly modest costs. This can be done by adding diverse altroutes, and by cross-connecting of routers and LANs. Many commands have multiple stovepipe networks that arrived from different programs and only need to be cross-connected. As DMS-like and NES-like security products enter the market and remove the reasons for segregating networks, or at least segregating WANs, this cross-connecting job gets easier.

Now, step back a bit and think about what the numbers mean. While the example component Ao values are fictitious, they are probably not too far from the truth. A 0.98 availability value may be acceptable for a garden variety office automation network, but won't be the instant you have some mission critical (C3I or combat system) functions floating around on it. And it may very well not be acceptable in the non-operational environment as soon as it gets popular with the boss. As soon as you eliminate all the single points of failure, you bump the Ao figures up to about 4 or 5 nines. Which is pretty respectable for C3I systems. Note that there's a pretty sharp knee in the curve between 2 and 4 nines. So don't quibble about values in between -- simply plan on dual-threading the system.[6]

If you look at the problem from a logistics and repairability point of view[7], this reinforces: even if the tech is standing by with the correct repair part in hand, nobody can repair equipment in less than half an hour.

You're forced to a dual-threaded system even if you work your numbers over a year instead of a month and then realize that you're likely to get something equivalent to one router hardware failure in that time -- and buy all of the downtime at once.

Note that there are major logistics savings in dual-threading systems as well. If one of the parallel components in our illustrated systems fails, the overall system continues to operate. And you can usually live with a 'next working day' repair regime which is much cheaper than '24 hour on-call'. Indeed, in shipboard environments, you have the flexibility to shift from organizational level

(ship's force) repair to dockside, depot level, repair which is usually a lot cheaper when you figure the costs of keeping qualified personnel and spare parts aboard. Since the vast majority of network components today are COTS, installing one or more spares in the network infrastructure is a fairly small capitalization cost.

Survivability. Once we understand the concepts of availability, thinking in terms of survivability adds only some minor twists.

The same principles apply, the only real differences are that instead of components frying themselves there is someone external who's trying to fry them. The same tools and arithmetic apply in analyzing survivability situations and designing to account for them.

As the Internet grows, there is a second twist, most noticable when examining security issues -- the good guys and the bad guys are on the same network. This provides some deterrent for the bad guys taking down the network -- it harms them too. This phenomenon is very pronounced in radionavigation systems -- moreso than in networks.

Conclusion. Requirements analysis of networks can resemble trying to ascend the down escalator if you attack the problem from a capacity analysis point of view. Whatever capacity requirements you calculate, they won't be the same when the network is actually installed. And there's a version of the Heisenberg Uncertainty Principle at work here too -- trying to estimate data rates will influence the estimates.

A far better approach is to examine the availability requirements first. Many times you will find that if you address the availability needs, you've taken care of the

capacity requirements in the wash. Or placed your network architecture in a position where capacity can easily be expanded later. Start here.

APPENDIX B

This appendix contains an excerpt from a thesis by Saravanan Radhakrishnan, Anupama Sundaresan, Gowri Dhandapani, written at the University of Kansas Information and Telecommunications Technology Center Department of Electrical Engineering & Computer Science, 19 December 1999. This brief discussion is meant to give the reader an overview of how differential service may be applied within a network.

The traffic classification process begins at the edge router or firewall. The edge router could be responsible for altering the TOS octet based upon the user profile originating the message. The edge router then places the outbound transmissions in a cue based on their traffic classification and employing the first in first out methodology then empties the cue. The core routers then have to only differentiate between the three levels of classification for transmission. This is simple but will provide a means of performing this task in a mission oriented method.

Present implementations are concerned largely with the prioritization of traffic based upon the type of application originating the traffic and not the source or individual or mission criticality behind the application [18]. The emphasis is on the quality of the transmission as received by the destination and not the value of the information within the transmission when compared to others. Tying the classification to the user profile creates a traceable link of responsibility to the individual originating the transmission. If greater

control over the use of the prioritization scheme is desired then the edge router could be used to check the destination IP address.

While it would be impractical to screen individual IP addresses, a simple check to see if the assigned IP address lies within the NMCI domain should be easy to accomplish. This would ensure that the services were being used for the appropriate purpose. This also gives the LAN administrator a mechanism for authenticating the originator at the edge router or the host, depending upon their desires, requirements, or configuration. In the end this employment of a differentiated services model could provide a means of providing a quality of service that would segregate traffic on a mission essential basis.

LIST OF REFERENCES

- [DONIT02] Department of the Navy IM/IT web page
<http://www.don-imit.navy.mil/interestTemplate.asp?type=initiative&theID=2&catID=1>
- [DAN00] Report to Congress on NMCI by the Secretary of the Navy, the Honorable Richard Danzig, of 20 May 2000, page A-1
- [CIP01] Navy and Marine Corps Intranet Brief presented to NMCI Information Bureau Oversight Council, 18 April 2001, Mr. Joseph Cipriano, Program Executive Officer for Information Technology, slide 3
- [WRD 98] Wired Magazine, 6.11 November 1998,
<http://www.wired.com/wired/archive//6.11/metcalfe.html>
- [DAN00] Executive Summary of the Report to Congress on NMCI by the Secretary of the Navy Richard Danzig, 20 May 2000.
- [RC02] Mr. Richard Clarke, Lecture 18 January 2002, NPS Monterey, CA
- [GCN02] Government Computer News, January 7, 2002; Vol. 21 No. 1, Dawn S. Olney
- [RJE99] *Requirements Definitions for Survivable Network Systems*: Carnegie Mellon University/Software Engineering Institute, May 1999, R. J Ellison, et al, p 1
- [RJE98] *A Case Study in Survivable Network System Analysis*, September 1998, R. J Ellison, et al, p 1
- [RE99] *Survivable Network Systems: An Emerging Discipline*, Carnegie Mellon University/Software Engineering Institute, May 1999, R. J Ellison, et al, p 5

- [NRM00] Survivable Network Analysis Method, Carnegie Mellon University/Software Engineering Institute, September 2000, N.R. Mead, et al, p5
- [RAY01] System Security Authorization Agreement (SSAA) for Navy/Marine Corps Intranet Raytheon Corporation dated 19 March 2001, p 8
- [RB 02] Professor Rex Buddenberg, Lecture notes of April 2002
<http://web1.nps.navy.mil/~budden/lecture.notes/availability.html>
- [MTX02] MATRIX ratings web site
<http://ratings.miq.net/compare-by-WREACHABILITY-US.html>
- [PEOIT02] NMCI Contract N00024-00-D-6000, Awarded 6 October 2000, pp 88 thru 96
- [CSC002] CISCO web site
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito doc/qos.htm#xtocid2>
- [JH97] *Net Gain, expanding markets through virtual communities*, John Hagel III, Authur G Armstrong, Harvard Business School Press, 1997
- [SBC02] SBC Communications website
<http://www.sbc.com/press room-/1,5932,31,00.html?query=20020507-1>
- [RC98] *Key Issues in Electronic Commerce and Electronic Publishing*, Roger Clarke, sections 4.1 and 4.2 <http://www.anu.edu.au/people/-Roger.Clarke/EC/Issues98.html#Iss2>
- [CERT02] *Overview of Attack Trends*, CERT Coordination Center, 8 April 2002, p 3
- [OPNAV83] OPNAVINST 4614.1F of 15 April 93, with changes 1 and 2.

- [SRSD99] *Diffspec - A Differentiated Services Tool*,
Saravanan Radhakrishnan, Anupama Sundaresan,
Gowri Dhandapani, Information and
Telecommunications Technology Center Department
of Electrical Engineering & Computer Science,
University of Kansas, 19 December 1999,
<http://qos.ittc.ukans.edu/DiffSpec/index.html>
- [RM02] ADM Richard Munns, June 2002, Senior Executive
Session, Naval Post Graduate School, Monterey,
CA
- [GCN02] Government Computer News, Volume 21 Number 9,
dated 29 April 2002, page 33
- [TR02] Tech Republic newsletter,
<http://www.techrepublic.com/article.jhtml?id=r00220020530hin01.htm&fromtm=e102-3>
- [LTG01] Navy Marine Corps Legacy Application Transition
Guide, version 2.1, dated 26 October 2001, p 3
- [PEOIT01] NMCI Update Briefing, NMCI Software Application
Management - PEO IT dtd 12 April 2001
- [CISCO02] CISCO web site,
<http://www.cisco.com/univercd/cc/td/doc/cisintw/k/ito doc/qos.htm#xtocid2>
- [TT02] Techtarget web site
http://searchSecurity.techtargget.com/sDefinition/0,,sid14_gci295031,00.html
- [RB97] Dr Rex Buddenberg, lectures notes of September
1997,
<http://web1.nps.navy.mil/~budden/lecture.notes/mahan.html>
- [BH01] *Active Defense; A comprehensive guide to
Network Security*, Chris Brenton and Cameron
Hunt, SYBEX Incorporated, Alameda CA

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, VA
2. Dudley Knox Library
Naval Postgraduate School
Monterey, CA
3. Dan C. Boger
Chairman, Information Sciences Dept
Naval Postgraduate School
Monterey, CA
4. Dr John Arquilla
Naval Postgraduate School
Monterey, CA
5. Professor Rex Buddenberg
Naval Postgraduate School
Monterey, CA
6. SECNAV DON CIO
1000 Navy Pentagon
Washington, DC 20350-1000
7. CNO N6
1000 Navy Pentagon
Washington, DC 20350-1000
8. Director NMCI Program Office
2231 Crystal City Drive
Suite 400
Arlington, VA 22202-3721
9. Joint Staff VJ6
ATTN ADM Brown
Pentagon
Washington, D.C. 20318-6000
10. Deputy PEO (USMC), Technical Director
Program Executive Office Information Technology
(PEO-IT)
2451 Crystal Drive, Suite 1109
Arlington, VA 22202-4804

11. Deputy PEO (Navy), Director NMCI Intranet Services
Program Executive Office Information Technology
(PEO-IT)
2451 Crystal Drive, Suite 1109
Arlington, VA 22202-4804
12. CERT® Coordination Center
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
13. COMNAVRESFOR N6
1000 Navy Pentagon
Washington, DC 20350-1000